

SHO-FA: Robust compressive sensing with order-optimal complexity, measurements, and bits

Mayank Bakshi, Sidharth Jaggi, Sheng Cai, Minghua Chen

{mayank, jaggi, cs010, minghua}@ie.cuhk.edu.hk

The Chinese University of Hong Kong, Hong Kong SAR China

Abstract

Suppose \mathbf{x} is any exactly k -sparse vector in \mathbb{R}^n . We present a class of “sparse” matrices A , and a corresponding algorithm that we call SHO-FA (for Short and Fast¹) that, with high probability over A , can reconstruct \mathbf{x} from $A\mathbf{x}$. The SHO-FA algorithm is related to the Invertible Bloom Lookup Tables (IBLTs) recently introduced by Goodrich *et al.*, with two important distinctions – SHO-FA relies on linear measurements, and is robust to noise. The SHO-FA algorithm is the first to simultaneously have the following properties: (a) it requires only $\mathcal{O}(k)$ measurements, (b) the bit-precision of each measurement and each arithmetic operation is $\mathcal{O}(\log(n) + P)$ (here 2^{-P} corresponds to the desired relative error in the reconstruction of \mathbf{x}), (c) the computational complexity of decoding is $\mathcal{O}(k)$ arithmetic operations, and (d) if the reconstruction goal is simply to recover a single component of \mathbf{x} instead of all of \mathbf{x} , with high probability over A this can be done in constant time. All constants above are independent of all problem parameters other than the desired probability of success. For a wide range of parameters these properties are information-theoretically order-optimal. In addition, our SHO-FA algorithm is robust to random noise, and (random) approximate sparsity for a large range of k . In particular, suppose the measured vector equals $A(\mathbf{x} + \mathbf{z}) + \mathbf{e}$, where \mathbf{z} and \mathbf{e} correspond respectively to the *source tail* and *measurement noise*. Under reasonable statistical assumptions on \mathbf{z} and \mathbf{e} our decoding algorithm reconstructs \mathbf{x} with an estimation error of $\mathcal{O}(\|\mathbf{z}\|_1 + (\log k)^2 \|\mathbf{e}\|_1)$. The SHO-FA algorithm works with high probability over A , \mathbf{z} , and \mathbf{e} , and still requires only $\mathcal{O}(k)$ steps and $\mathcal{O}(k)$ measurements over $\mathcal{O}(\log(n))$ -bit numbers. This is in contrast to most existing algorithms which focus on the “worst-case” \mathbf{z} model, where it is known $\Omega(k \log(n/k))$ measurements over $\mathcal{O}(\log(n))$ -bit numbers are necessary. Our algorithm has good empirical performance, as validated by simulations.

I. INTRODUCTION

In recent years, spurred by the seminal work on *compressive sensing* of [1], [2], much attention has focused on the problem of reconstructing a length- n “compressible” vector \mathbf{x} over \mathbb{R} with fewer than n linear measurements. In particular, it is known (e.g. [3], [4]) that with $m = \mathcal{O}(k \log(n/k))$ linear measurements one can computationally efficiently² obtain a vector $\hat{\mathbf{x}}$ such that the *reconstruction error* $\|\mathbf{x} - \hat{\mathbf{x}}\|_1$ is $\mathcal{O}(\|\mathbf{x} - \mathbf{x}_k^*\|_1)$,³ where \mathbf{x}_k^* is the best possible k -sparse approximation to \mathbf{x} (specifically, the k non-zero terms of \mathbf{x}_k^* correspond to the k largest components of \mathbf{x} in magnitude, hence $\mathbf{x} - \mathbf{x}_k^*$ corresponds to the “tail” of \mathbf{x}). A number of different classes of algorithms are able to give such performance, such as those based on ℓ_1 -optimization (e.g. [1], [2]), and those based on iterative “matching pursuit” (e.g. [8], [9]). Similar results, with an additional additive term in the reconstruction error hold even if the linear measurements themselves also have noise added to them (e.g. [3], [4]). The fastest of these algorithms use ideas from the theory of expander graphs, and have running time $\mathcal{O}(n \log(n/k))$ [10]–[12].

The class of results summarized above are indeed very strong – they hold for *all* \mathbf{x} vectors, including those with “worst-case tails”, *i.e.* even vectors where the components of \mathbf{x} smaller than the k largest coefficients (which can be thought of as “source tail”) are chosen in a maximally worst-case manner. In fact [13] proves that to obtain a reconstruction error that scales linearly with the ℓ_1 -norm of the \mathbf{z} , the tail of \mathbf{x} , requires $\Omega(k \log(n/k))$ linear measurements.

¹Also, SHO-FA sho good! In fact, it’s all $\mathcal{O}(k)$!

²The caveat is that the reconstruction techniques require one to solve an LP. Though polynomial-time algorithms to solve LPs are known, they are generally considered to be impractical for large problem instances.

³In fact this is the so-called $\ell_1 < C\ell_1$ guarantee. One can also prove stronger $\ell_2 < C\ell_1/\sqrt{k}$ reconstruction guarantees for algorithms with similar computational performance, and it is known that a $\ell_2 < C\ell_2$ reconstruction guarantee is not possible if the algorithm is required to be zero-error [5], but is possible if some (small) probability of error is allowed [6], [7].

Number of measurements: However, depending on the application, such a lower bound based on “worst-case \mathbf{z} ” may be unduly pessimistic. For instance, it is known that if \mathbf{x} is exactly k -sparse (has exactly k non-zero components, and hence $\mathbf{z} = \mathbf{0}$), then based on Reed-Solomon codes [14] one can efficiently reconstruct \mathbf{x} with $\mathcal{O}(k)$ noiseless measurements (e.g. [15]) via algorithms with decoding time-complexity $\mathcal{O}(n \log(n))$, or via codes such as in [16], [17] with $\mathcal{O}(k)$ noiseless measurements with decoding time-complexity $\mathcal{O}(n)$.⁴ In the regime where $k = \theta(n)$ [18] use the “sparse-matrix” techniques of [10]–[12] to demonstrate that $\mathcal{O}(k) = \mathcal{O}(n)$ measurements suffice to reconstruct \mathbf{x} .

Noise: Even if the source is not exactly k -sparse, a spate of recent work has taken a more information-theoretic view than the coding-theoretic/worst-case point-of-view espoused by much of the compressive sensing work thus far. Specifically, suppose the length- n source vector is the sum of *any* exactly k -sparse vector \mathbf{x} and a “random” source noise vector \mathbf{z} (and possibly the linear measurement vector $A(\mathbf{x} + \mathbf{z})$ also has a “random” measurement noise vector \mathbf{e} added to it). Then as long as the noise variances are not “too much larger” than the signal power, the work of [19] demonstrates that $\mathcal{O}(k)$ measurements suffice (though the proofs in [19] are information-theoretic and existential – the corresponding “typical-set decoding” algorithms require time exponential in n). Indeed, even the work of [13], whose primary focus was to prove that $\Omega(k \log(n/k))$ linear measurements are necessary to reconstruct \mathbf{x} in the worst case, also notes as an aside that if \mathbf{x} corresponds to an exactly sparse vector plus random noise, then in fact $\mathcal{O}(k)$ measurements suffice. The work in [20], [21] examines this phenomenon information-theoretically by drawing a nice connection with the *Rényi information dimension* $\bar{d}(X)$ of the signal/noise, and [22] show how to computationally efficiently achieve this performance by exactly reconstructing \mathbf{x} with $\mathcal{O}(\bar{d}(X)n) + o(n)$ samples in time $\mathcal{O}(n)$. Corresponding lower bounds showing $\Omega(k \log(n/k))$ samples are required in the higher noise regime are provided in [23], [24].

Number of measurement bits: However, most of the works above focus on minimizing the number of linear measurements in $A\mathbf{x}$, rather than the more information-theoretic view of trying to minimize the number of bits in $A\mathbf{x}$ over all measurements. Some recent work attempts to fill this gap – notably “Counting Braids” [25], [26] (this work uses “multi-layered non-linear measurements”), and “one-bit compressive sensing” [27], [28] (the corresponding decoding complexity is somewhat high (though still polynomial-time) since it involves solving an LP).

Decoding time-complexity: The emphasis of the discussion thus far has been on the number of linear measurements/bits required to reconstruct \mathbf{x} . The decoding algorithms in most of the works above have decoding time-complexities⁵ that scale at least linearly with n . In regimes where k is significantly smaller than n , it is natural to wonder whether one can do better. Indeed, algorithms based on iterative techniques answer this in the affirmative. These include Chaining Pursuit [29], group-testing based algorithms [30], and Sudocodes [31] – each of these have decoding time-complexity that can be sub-linear in n (but at least $\mathcal{O}(k \log(k) \log(n))$), but each requires at least $\mathcal{O}(k \log(n))$ linear measurements.

Database query: Finally, we consider a *database query* property that is not often of primary concern in the compressive sensing literature. That is, suppose one is given a compressive sensing algorithm that is capable of reconstructing \mathbf{x} with the desired reconstruction guarantee. Now suppose that one instead wishes to reconstruct, with reasonably high probability, just “a few” (constant number) *specific* components of \mathbf{x} , rather than all of it. Is it possible to do so even faster (say in constant time) – for instance, if the measurements are in a database, and one wishes to query it in a computationally efficient manner? If the matrix A is “dense” (most of its entries are non-zero) then one can directly see that this is impossible. However, several compressive sensing algorithms (for instance [18]) are based on “sparse” matrices A , and it can be shown that in fact these algorithms do indeed have this property “for free” (as indeed does our algorithm), even though the authors do not analyze this. As can be inferred from the name, this database query property is more often considered in the database community, for instance in the work on IBLTs [32].

⁴In general the linear systems produced by Reed-Solomon codes are ill-conditioned, which causes problems for large n .

⁵For ease of presentation, in accordance with common practice in the literature, in this discussion we assume that the time-complexity of performing a single arithmetic operation is constant. Explicitly taking the complexity of performing finite-precision arithmetic into account adds a multiplicative factor (corresponding to the precision with which arithmetic operations are performed) in the time-complexity of most of the works, including ours.

A. Our contributions

Conceptually, the “iterative decoding” technique we use is not new. Similar ideas have been used in various settings in, for instance [16], [32]–[34]. However, to the best of our knowledge, no prior work has the same performance as our work – namely – information-theoretically order-optimal number of measurements, bits in those measurements, and time-complexity, for the problem of reconstructing a sparse signal (or sparse signal with a noisy tail and noisy measurements) via linear measurements (along with the database query property).⁶ The key to this performance is our novel design of “sparse random” linear measurements, as described in Section II.

To summarize, the desirable properties of SHO-FA are that with high probability⁷:

- **Number of measurements:** For every k -sparse \mathbf{x} , with high probability over A , $\mathcal{O}(k)$ linear measurements suffice to reconstruct \mathbf{x} . This is information-theoretically order-optimal.
- **Number of measurement bits:** The total number of bits in $A\mathbf{x}$ required to reconstruct \mathbf{x} to a relative error of 2^{-P} is $\mathcal{O}(k(\log(n) + P))$. This is information-theoretically order-optimal for any $k = \mathcal{O}(n^{1-\Delta})$ (for any $\Delta > 0$).
- **Decoding time-complexity:** The total number of arithmetic operations required is $\mathcal{O}(k)$. This is information-theoretically order-optimal.
- **Database queries:** With constant probability $1 - \epsilon$ any single database query can be answered in constant time.⁸
- **Noise:** Suppose \mathbf{z} and \mathbf{e} have i.i.d. components⁹ drawn respectively from $\mathcal{N}(0, \sigma_z^2)$ and $\mathcal{N}(0, \sigma_e^2)$. For $k = \mathcal{O}(n^{1-\Delta})$ for any $\Delta > 0$, a modified version of SHO-FA (mod-SHO-FA) that with high probability reconstructs \mathbf{x} with an estimation error of $\mathcal{O}(\|\mathbf{z}\|_1 + (\log k)^2 \|\mathbf{e}\|_1)$.
- **Practicality:** As validated by simulations (shown in Appendix I), most of the constant factors involved above are not large, and are in fact significantly smaller than the explicit constants that can be calculated via our analysis.
- **Different bases:** As is common in the compressive sensing literature, our techniques generalize directly to the setting wherein \mathbf{x} is sparse in an alternative basis (say, for example, in a wavelet basis). We defer discussion of this until after the description of SHO-FA.
- **Universality:** While we present a specific ensemble of matrices over which SHO-FA operates, we argue that in fact similar algorithms work over fairly general ensembles of “sparse random matrices”, and further that such matrices can occur in applications, for instance in wireless MIMO systems [38]. Again, we defer discussion of this issue to after our description of SHO-FA.

B. Special acknowledgements

In particular, the bounds on the minimum number of measurements required for “worst-case” recovery and the corresponding discussion on recovery of signals with “random tails” in [13] led us to consider this problem in the first place. Equally, the class of compressive sensing codes in [18], which in turn build upon the constructions of expander codes in [33], have been influential in leading us to this work. While the model in [34] differs from the one in this work, the techniques therein are of significant interest in our work. The analysis in [34] of the number of disjoint components in certain classes of random graphs, and also the analysis of how noise propagates in iterative decoding is potentially useful sharpening our results. We elaborate on these in Section III.

⁶While writing this paper, we became aware of a parallel work by Pawar and Ramchandran [35] that seems to rely on ideas similar to our work and achieves similar performance guarantees. However, at the time of submission, a preprint of the work by Pawar and Ramchandran is not available for us to compare and contrast the two results.

⁷For most of the properties, we show that this probability is at least $1 - 1/k^{\mathcal{O}(1)}$, though we explicitly prove only $1 - \mathcal{O}(1/k)$.

⁸The constant ϵ can be made arbitrarily close to zero, at the cost of a multiplicative factor $\mathcal{O}(1/\epsilon)$ in the number of measurements required. In fact, if we allow the number of measurements to scale as $\mathcal{O}(k \log(k))$, we can support any number of database queries, each in constant time, with probability of every one being answered correctly at with probability at least $1 - \epsilon$.

⁹Even if the statistical distribution of the components of \mathbf{z} and \mathbf{e} are not i.i.d. Gaussian, statements with a similar flavor can be made. For instance, pertaining to the effect of the distribution of \mathbf{z} , it turns out that our analysis is sensitive only on the distribution of the *sum* of components of \mathbf{z} , rather than the components themselves. Hence, for example, if the components of \mathbf{z} are i.i.d. non-Gaussian, it turns out that via the Berry-Esseen theorem [36] one can derive similar results to the ones derived in this work. In another direction, if the components of \mathbf{z} are not i.i.d. but do satisfy some “regularity constraints”, then using Bernstein’s inequality [37] one can again derive analogous results. However, these arguments are more sensitive and outside the scope of this paper, where the focus is on simpler models.

Reference	A	\mathbf{x}	\mathbf{z}	\mathbf{e}	Reconstruction Goal	\mathbf{P}_e	# Measurements	# Decoding steps	Precision
Reed-Solomon [14]	D	D	0	0	Exact	0	$2k + 1$	$\mathcal{O}(n \log n)$ [39]	–
Singleton [40]	D/R	D	0	0	Exact	0	$\geq 2k$	–	–
Mitzenmacher-Varghese [17]	R	D	0	0	Exact	$\mathcal{O}(1)$	$\mathcal{O}(k)$	$\mathcal{O}(n)$	–
Kudekar-Pfister [16]	R	D	0	0	Exact	$o(1)$	$\mathcal{O}(k)$	$\mathcal{O}(n)$	–
Tropp-Gilbert [8]	G	D	0	0	Exact	$o(1)$	$\mathcal{O}(k \log(n))$	$\mathcal{O}(k^2 n \log(n))$	–
Wu-Verdú '10 [20]	R	R	R	0	Exact	$o(1)$	$nd(\mathbf{x} + \mathbf{z}) + o(n)$	$\mathcal{O}(\exp(n))$	–
Donoho <i>et al.</i> [22]	R	R	R	0 R	Exact $l_2 < Cl_2$	$o(1)$	$nd(\mathbf{x} + \mathbf{z}) + o(n)$	$\mathcal{O}(n^3)$	–
Cormode-Muthukrishnan [30]	R	D	0	0	$l_2 < Cl_2$	$o(1)$	$\mathcal{O}(k \text{ poly}(\log(n)))$	$\mathcal{O}(k \text{ poly}(\log(n)))$	–
Cohen <i>et al.</i> [5]	D	D	D	0	$l_2 < Cl_2$	0	$\Omega(n)$	–	–
Price-Woodruff [7]	D	D	D	0	$l_2 < Cl_2$	$o(1)$	$\theta(k \log(n/k))$	–	–
Ba <i>et al.</i> [13]	D/R	D	D	0	$l_1 < Cl_1$	$\mathcal{O}(1)$	$\Omega(k \log(n/k))$	–	$\mathcal{O}(\log(n))$
Ba <i>et al.</i> [13]	R	D	R	0	$l_2 < Cl_2$	$o(1)$	$\mathcal{O}(k)$	$\mathcal{O}(\exp(n))$	–
Candés [3], Baraniuk <i>et al.</i> [4]	R	D	D	D	$l_2 < \frac{C}{\sqrt{k}} l_1$	$o(1)$	$\mathcal{O}(k \log(n/k))$	LP	–
Indyk <i>et al.</i> [41]	D	D	D	D	$l_1 < (1 + \epsilon) l_1$	0	$\mathcal{O}(k \log(n/k))$	$\mathcal{O}(n \log(n/k))$	–
Akçakaya <i>et al.</i> [42]	R	D	0	R	$l_2 < Cl_2$ / Sup. Rec.	0	$\mathcal{O}(k)$ Cond. on x_{\min}	$\mathcal{O}(\exp(n))$	–
Wu-Verdú '11 [21]	R	R	R	R	$l_2 < Cl_2$	$\mathcal{O}(1)$	$d(\mathbf{x} + \mathbf{z})$	$\mathcal{O}(\exp(n))$	–
Wainwright [24]	\mathcal{N}	D	0	R	Sup. Rec.	$\mathcal{O}(1)$	$\Omega(k \log(n/k))$	–	–
Fletcher <i>et al.</i> [23]	\mathcal{N}	D	0	R	Sup. Rec.	$o(1)$	$\mathcal{O}(k \log(n - k))$	–	–
Aeron <i>et al.</i> [43]	\mathcal{N}	D	0	R	Sup. Rec.	$\mathcal{O}(1)$	$\Omega(k \log(n/k))$	–	–
Plan-Vershynin [27]	R	D	0	sgn	$l_2 < f(\mathbf{x}, \mathbf{x}_k)$	$\mathcal{O}(1)$	$k^2 \log(n/k)$	LP	1
Jacques <i>et al.</i> [28]	R	D	0	sgn	$l_2 < f'(n, k)$	$\mathcal{O}(1)$	$k \log(n)$	$\exp(n)$	1
Sarvotham <i>et al.</i> [31]	R	D	0	0	Exact	$\mathcal{O}(1)$	$k \log(n)$	$k \log(k) \log(n)$	–
Gilbert <i>et al.</i> [29]	R	P.L.	P.L.	0	$l_1 < [1 + C \log(n)] l_1$	$\mathcal{O}(1)$	$k \log^2(n)$	$k \log^2(n) \log^2(k)$	–
This work	R R	D D	0 R	0 R	Exact $l_1 < Cl_1$	$o(1)$ $o(1)$	$\mathcal{O}(k)$ $\mathcal{O}(k)$	$\mathcal{O}(k)$ $\mathcal{O}(k)$	$\mathcal{O}(\log(n) + P)$ $\mathcal{O}(\log(n))$

Explanations and discussion: At the risk of missing much of the literature, and also perhaps oversimplifying nuanced results, we summarize in this table many of the strands of work preceding this paper and related to it – not all results from each work are represented in this table. The second to the fifth columns respectively reference whether the measurement matrix A , source k -sparse vector \mathbf{x} , source noise \mathbf{z} , and measurement noise \mathbf{e} are random (R) or deterministic (D) – a 0 in a column corresponding to noise indicates that that work did not consider that type of noise. An entry “P.L.” stands for “Power Law” decay in columns corresponding to \mathbf{x} and \mathbf{z} . For achievability schemes, in general D -type results are stronger than R -type results, which in turn are stronger than 0-type results. This is because a D -type result for the measurement matrix indicates that there is an explicit construction of a matrix that satisfies the required goals, whereas the R -type results generally indicate that the result is true with high probability over measurement matrices. Analogously, a D in the columns corresponding to \mathbf{x} , \mathbf{z} or \mathbf{e} indicates that the scheme is true for all vectors, whereas an R indicates that it is true for random vectors from some suitable ensemble. For converse results, the the opposite is true 0 results are stronger than R -type results, which are stronger than D -type results. An entry \mathcal{N} indicates the normal distribution – the results of [24] and [23] are converses for matrices with i.i.d. Gaussian entries. An entry “sgn” denotes (in the case of works dealing with one-bit measurements) that the errors are sign errors. The sixth column corresponds to what the desired goal is. The strongest possible goal is to have exact reconstruction of \mathbf{x} (up to quantization error due to finite-precision arithmetic), but this is not always possible, especially in the presence of noise. Other possible goals include “Sup. Rec.” (short for support recovery) of \mathbf{x} , or that the reconstruction $\hat{\mathbf{x}}$ of \mathbf{x} differs from \mathbf{x} as a “small” function of \mathbf{z} . It is known that if a deterministic reconstruction algorithm is desired to work for all \mathbf{x} and \mathbf{z} , then $\|\hat{\mathbf{x}} - \mathbf{x}\|_2 < \mathcal{O}(\|\mathbf{z}\|_2)$ is not possible with less than $\Omega(n)$ measurements [5], and that $\|\hat{\mathbf{x}} - \mathbf{x}\|_2 < \mathcal{O}(\|\mathbf{z}\|_1/\sqrt{k})$ implies $\|\hat{\mathbf{x}} - \mathbf{x}\|_1 < \mathcal{O}(\|\mathbf{z}\|_1)$. The reconstruction guarantees in [27], [28] unfortunately do not fall neatly in these categories. The seventh column indicates what the probability of error is – *i.e.* the probability over any randomness in A , \mathbf{x} , \mathbf{z} and \mathbf{e} that the reconstruction goal in the sixth column is not met. In the eighth column, some entries are marked $d(\mathbf{x} + \mathbf{z})$ – this denotes the (upper) Rényi dimension of $\mathbf{x} + \mathbf{z}$ – in the case of exactly k -sparse vectors this equals k , but for non-zero \mathbf{z} it depends on the distribution of \mathbf{z} . The ninth column considers the computational complexity of the algorithms – the entry “LP” denotes the computational complexity of solving a linear program. The final column notes whether the particular work referenced considers the precision of arithmetic operations, and if so, to what level.

The work that is conceptually the closest to SHO-FA is that of the Invertible Bloom Lookup Tables (IBLTs) introduced by Goodrich-Mitzenmacher [32] (though our results were derived independently, and hence much of our analysis follows a different line of reasoning). The data structures and iterative decoding procedure (called “peeling” in [32]) used are structurally very similar to the ones used in this work. However the “measurements” in IBLTs are fundamentally non-linear in nature – specifically, each measurement includes within it a “counter” variable – it is not obvious how to implement this in a linear manner. Therefore, though the underlying graphical structure of our algorithms is similar, the details of our implementation require new non-trivial ideas. Also, IBLTs as described are not robust to either signal tails or measurement noise. Nonetheless, the ideas in [32] have been influential in this work. In particular, the notion that an individual component of \mathbf{x} could be recovered in constant

time, a common feature of Bloom filters, came to our notice due to this work. Also, the analysis (via “2-cores of random hypergraphs”) in [32] of the leading constant factor in the number of measurements m is tighter than that presented in our work. Since the emphasis of our work is on order optimality, we choose to present our (fairly straightforward) analysis, and leave for future work the incorporation of their more technical line of reasoning into our algorithm.

II. EXACTLY k -SPARSE \mathbf{x} AND NOISELESS MEASUREMENTS

We first consider the simpler case when the source signal is exactly k -sparse and the measurements are noiseless, *i.e.*, $\mathbf{y} = A\mathbf{x}$, and both \mathbf{z} and \mathbf{e} are all-zero vectors. The intuition presented here carries over to the scenario wherein both \mathbf{z} and \mathbf{e} are non-zero, considered separately in Section III

For k -sparse input vectors $\mathbf{x} \in \mathbb{R}^n$ let the set $\mathcal{S}(\mathbf{x})$ denote its *support*, *i.e.*, its set of nonzero values $\{j : x_j \neq 0\}$. Recall that in our notation, for some m , a *measurement matrix* $A \in \mathbb{R}^{m \times n}$ is chosen probabilistically. This matrix operates on \mathbf{x} to yield the *measurement vector* $\mathbf{y} \in \mathbb{R}^m$ as $\mathbf{y} = A\mathbf{x}$. The decoder takes the vector \mathbf{y} as input and outputs the reconstruction $\hat{\mathbf{x}} \in \mathbb{R}^n$ – it is desired that $\hat{\mathbf{x}}$ equal \mathbf{x} (with upto P bits of precision) with high probability over the choice of measurement matrices.

In this section, we describe a probabilistic construction of the measurement matrix A and a reconstruction algorithm SHO-FA that achieves the following guarantees.

Theorem 1. *Let $k \leq n$. There exists a reconstruction algorithm SHO-FA for $A \in \mathbb{R}^{m \times n}$ with the following properties:*

- 1) *For every $\mathbf{x} \in \mathbb{R}^n$, with probability $1 - \mathcal{O}(1/k)$ over the choice of A , SHO-FA produces a reconstruction $\hat{\mathbf{x}}$ such that $\|\mathbf{x} - \hat{\mathbf{x}}\|_1 / \|\mathbf{x}\|_1 \leq 2^{-P}$*
- 2) *$m = ck$ for some $c > 0$*
- 3) *Expected number of steps required by SHO-FA is $\mathcal{O}(k)$*
- 4) *Expected number of bitwise arithmetic operations required by SHO-FA is $\mathcal{O}(k(\log n + P))$.*

We prove the above theorem in the remainder of this section.

A. High-level intuition

If $m = \Theta(n)$, the task of reconstructing \mathbf{x} from $\mathbf{y} = A\mathbf{x}$ appears similar to that of *syndrome decoding* of a channel code of rate n/m [44]. It is well-known [45] that channel codes based on *bipartite expander graphs*, *i.e.*, bipartite graphs with good expansion guarantees for all sets of size less than or equal to k , allow for decoding in a number of steps that is linear in the size of \mathbf{x} . In particular, given such a bipartite expander graph with n nodes on the left and m nodes on the right, choosing the matrix A as a $m \times n$ binary matrix with non-zero values in the locations where the corresponding pair of nodes in the graph has an edge is known to result in codes with rate and relative minimum distance that is linear in n .

Motivated by this [18] explore a measurement design that is derived from expander graphs and show that $\mathcal{O}(k \log(n/k))$ measurements suffice, and $\mathcal{O}(k)$ iterations with overall decoding complexity of $\mathcal{O}(n \log(n/k))$.¹⁰

It is tempting to think that perhaps an optimized application of expander graphs could result in a design that require only $\mathcal{O}(k)$ number of measurements. However, we show that in the compressive sensing setting, where, typically $k = o(n)$, it is not possible to satisfy the desired expansion properties. In particular, if one tries to mimic the approach of [18], one would need bipartite expanders such that *all* sets of size k on one side of the graph “expand” – we show in Lemma 2 that this is not possible. As such, this result may be of independent interest for other work that require similar graphical constructions (for instance the “magical graph” constructions of [46], or the expander code constructions of [33] in the high-rate regime).

Instead, one of our key ideas is that we do not really need “true” expansion. Instead, we rely on a notion of approximate expansion that guarantees expansion for most k -sized sets (and their subsets) of nodes on the left of our bipartite graph. We do so by showing that any set of size at most k , with high probability over suitably chosen

¹⁰The work of [10] is related – it also relies on bipartite expander graphs, and has similar performance for exactly k -sparse vectors. But [10] can also handle general approximately k -sparse vectors, unlike [18]. However, our algorithms are closer in spirit to those of [18], and hence we focus on this work.

measurement matrices, expands to the desired amount. Probabilistic constructions turn out to exist for our desired property.¹¹ Such a construction is shown in Lemma 1.

Our second key idea is that in order to be able to recover all the k non-zero components of \mathbf{x} with at most $\mathcal{O}(k)$ steps in the decoding algorithm, it is necessary (and sufficient) that on average, the decoder reconstructs one previously undecoded non-zero component of \mathbf{x} , say x_j , in $\mathcal{O}(1)$ steps in the decoding algorithm. For $k = o(n)$ the algorithm does not even have enough time to write out all of \mathbf{x} , but only its non-zero values. To achieve such efficient identification of x_j , we go beyond the 0/1 matrices used in almost all prior work on compressive sensing based on expander graphs.¹² Instead, we use distinct values in each row for the non-zero values in A , so that if only one non-zero x_j is involved in the linear measurement involving a particular y_i (a situation that we demonstrate happens in a constant fraction of y_i), one can identify which x_j it must be in $\mathcal{O}(1)$ time. Our decoding then proceeds iteratively, by identifying such x_j and canceling their effects on y_i , and terminates after $\mathcal{O}(k)$ steps after all non-zero x_j and their locations have been identified (since we require our algorithm to work with high probability for all \mathbf{x} , we also add “verification” measurements – this only increases the total number of measurements by a constant factor). Our calculations are precise to $\mathcal{O}(\log(n) + P)$ bits – the first term in this comes from requirements necessary for computationally efficient identification of non-zero x_j , and the last term from the requirement that we require that the reconstructed vector be correct up to P -precision. Hence the total number of bits over all measurements is $\mathcal{O}(k((\log(n) + P)))$. Note that this is information-theoretically order-optimal, since even specifying k locations in a length- n vector requires $\Omega(k(\log(n/k)))$ bits, and specifying the value of the non-zero locations so that the relative reconstruction error is $\mathcal{O}(2^{-P})$ requires $\Omega(kP)$ bits.

We now present our SHO-FA algorithm in two stages. We first by use our first key idea (of “approximate”) expansion in Section to describe some properties of bipartite expander graphs with certain parameters. We then show in Section how these properties, via our second key idea (of efficient identification) can be used by SHO-FA to obtain desirable performance.

B. Description of graph properties

We first construct a bipartite graph \mathcal{G} (see Example 1 in the Appendix) with some desirable properties outlined below. We then show in Lemmas 1 and 3 that such graphs exist (Lemma 2 shows the non-existence of graphs with even stronger properties). In Section II-C we then use these graph properties in the SHO-FA algorithm. To simplify notation in what follows (unless otherwise specified) we omit rounding numbers resulting from taking ratios or logarithms, with the understanding that the corresponding inaccuracy introduced is negligible compared to the result of the computation. Also, for ease of exposition, we fix various internal parameters to “reasonable” values rather than optimizing them to obtain “slightly” better performance at the cost of obfuscating the explanations – whenever this happens we shall point it out parenthetically. Lastly, let ϵ be any “small” positive number, corresponding to the probability of a certain “bad event”.

Properties of \mathcal{G} :

- 1) *Construction of a left-regular bipartite graph:* The graph \mathcal{G} is chosen uniformly at random from the set of bipartite graphs with n nodes on the left and m' nodes on the right, such that each node on the left has degree $d \geq 7$. In particular, m' is chosen to equal ck for some design parameter c to be specified later as part of code design.
- 2) *Edge weights for “identifiability”:* For each node on the right, the weights of the edges attached to it are required to be distinct. In particular, each edge weight is chosen as a complex number of unit magnitude, and phase between 0 and $\pi/2$. Since there are a total of dn edges in \mathcal{G} , choosing distinct phases for each edge attached to a node on the right requires at most $\log(dn)$ bits of precision (though on average there are about dn/m' edges attached to a node on the right, and hence on average one needs about $\log(dn/m')$ bits of precision).
- 3) *$\mathcal{S}(\mathbf{x})$ -expansion:* With high probability over \mathcal{G} defined in Property 1 above, for any set $\mathcal{S}(\mathbf{x})$ of k nodes on the left, the number of nodes neighbouring those in any $\mathcal{S}'(\mathbf{x}) \subseteq \mathcal{S}(\mathbf{x})$ is required to be at least $2d/3$ times

¹¹In fact similar properties have been considered before in the literature – for instance [46] constructed so-called “magical graphs” with similar properties. Our contribution is the way we use this property for our needs.

¹²It can be argued that such a choice is a historical artifact, since error-correcting codes based on expanders were originally designed to work over the binary field \mathbb{F}_2 . There is no reason to stick to this convention when, as now, computations are done over \mathbb{R} .

the size of $\mathcal{S}'(\mathbf{x})$.¹³ The proof of this statement is the subject of Lemma 1.

- 4) “Many” $\mathcal{S}(\mathbf{x})$ -leaf nodes: For any set $\mathcal{S}(\mathbf{x})$ of at most k nodes on the left of \mathcal{G} , we call any node on the right of \mathcal{G} an $\mathcal{S}(\mathbf{x})$ -leaf node if it has exactly one neighbor in $\mathcal{S}(\mathbf{x})$, and we call it a $\mathcal{S}(\mathbf{x})$ -non-leaf node if it has two or more neighbours in $\mathcal{S}(\mathbf{x})$. (If the node on the right has no neighbours in $\mathcal{S}(\mathbf{x})$, we call it a $\mathcal{S}(\mathbf{x})$ -zero node.) Assuming $\mathcal{S}(\mathbf{x})$ satisfies the expansion condition in Property 3 above, it can be shown that at least a fraction $1/2$ of the nodes that are neighbours of any $\mathcal{S}'(\mathbf{x}) \subseteq \mathcal{S}(\mathbf{x})$ are $\mathcal{S}'(\mathbf{x})$ -leaf nodes.¹⁴ The proof of this statement is the subject of Lemma 1.

Example 1: We now demonstrate via the following toy example in Figures II-B and 2 a graph \mathcal{G} satisfying Properties 1-4.




Figure 1. *Property 1:* Bipartite “approximate expander” graph with $n = 5$ nodes on the left, and $m' = 4$ nodes on the right. Each node on the left has degree 3. *Property 2:* The thicknesses of the edges represent the weights assigned to the edges. In particular, it is required that for each node on the right, the edges incoming have distinct weights. In this example, the thinnest edges are assigned a weight of 1, the next thickest edges have a weight $e^{i\pi/6}$, the next thickest edges have weight $e^{i2\pi/6} = e^{i\pi/3}$, and the thickest edges have weight $e^{i3\pi/6} = e^{i\pi/2}$.




Figure 2. *Property 3:* We require that *most* sets $\mathcal{S}'(\mathbf{x})$ of at most $|\mathcal{S}(\mathbf{x})| = k = 2$ nodes on the left in the graph \mathcal{G} in Figure II-B have at least $2|\mathcal{S}'(\mathbf{x})|$ neighbors on the right. In the graph in Figure II-B it can be manually verified that most sets of size $\mathcal{S}'(\mathbf{x})$ at most 2 have at least $2|\mathcal{S}'(\mathbf{x})|$ neighbors. For example, Figure II-B(a) focuses on the subset $\mathcal{S}'(\mathbf{x}) = \{1, 5\}$ of nodes on the left side of \mathcal{G} in Figure II-B. This particular $\mathcal{S}'(\mathbf{x})$ has 4 neighbours, and all its single-node subsets have 3 neighbours. The only $\mathcal{S}'(\mathbf{x})$ set of two or fewer nodes that does not satisfy Property 3 is $\{2, 5\}$, as shown in Figure II-B(b), since it has only $3 < 2 \times 2$ neighbours. *Property 4:* For sets $\mathcal{S}'(\mathbf{x})$ that satisfy Property 3 it can be manually verified that “many” of their neighbours are $\mathcal{S}'(\mathbf{x})$ -leaf nodes. For example, for $\mathcal{S}'(\mathbf{x}) = \{1, 5\}$, two out of its four neighbours (i.e., a fraction $1/2$) are $\mathcal{S}'(\mathbf{x})$ -leaf nodes – which satisfies the constraint that at least a fraction $1/2$ of its neighbours be $\mathcal{S}'(\mathbf{x})$ -leaf nodes. On the other hand, for $\mathcal{S}'(\mathbf{x}) = \{2, 5\}$ (which does *not* satisfy Property 3), *none* of its neighbours are $\mathcal{S}'(\mathbf{x})$ -leaf nodes.

□

We now state the Lemmas needed to make our arguments precise. First, we formalize the $\mathcal{S}'(\mathbf{x})$ -expansion property defined in Property 3.

Lemma 1. (Property 3 ($\mathcal{S}(\mathbf{x})$ -expansion)): *Let $\epsilon > 0$ and $k < n \in \mathbb{N}$ be arbitrary, and let $c \in \mathbb{N}$ be fixed. Let \mathcal{G} be chosen uniformly at random from the set of all bipartite graphs with n nodes (each of degree d) on the left*

¹³The expansion factor $2d/3$ is somewhat arbitrary. In our proofs, this can be replaced with any number strictly between half the degree and the degree of the left nodes, and indeed one can carefully optimize over such choices so as to improve the constant in front of the expected time-complexity/number of measurements of SHO-FA. Again, we omit this optimization since this can only improve the performance of SHO-FA by a constant factor.

¹⁴Yet again, this choice of $1/2$ is a function of the choices made for the degree of the left nodes in Property 1 and the expansion factor 2 in Property 3. Again, we omit optimizing it.

and $m' = ck$ nodes on the right. Then for any $S(\mathbf{x})$ of size at most k and any $S'(\mathbf{x}) \subseteq S(\mathbf{x})$, with probability $1 - o(1/k)$ (over the random choice \mathcal{G}) there are at least $2d/3$ times as many nodes neighbouring those in $S'(\mathbf{x})$, as there are in $S'(\mathbf{x})$.

Proof: Follows from a standard probabilistic method argument. Given for completeness in Appendix A.

Note here that, in contrast to the “usual” definition of “vertex expansion” [45] (wherein the expansion property is desired “for all” subsets of left nodes up to a certain size) Lemma 1 above only gives a probabilistic expansion guarantee for any subset of $S(\mathbf{x})$ of size k . In fact, Lemma 2 below shows that for the parameters of interest, “for all”-type expanders cannot exist.

Lemma 2. *Let $k = o(n)$, and $d > 0$ be an arbitrary constant. Let \mathcal{G} be an arbitrary bipartite graph with n nodes (each of degree d) on the left and m' nodes on the right. Then for all sufficiently large n , suppose each set of size k of $S(\mathbf{x})$ nodes on the left of \mathcal{G} has strictly more than $d/2$ times as many nodes neighbouring those in $S(\mathbf{x})$, as there are in $S(\mathbf{x})$. Then $m' = \Omega(k \log(n/k))$.*

Proof: Follows from the Hamming bound in coding theory [44] and standard techniques for expander codes [33]. Proof in Appendix B.

Another way of thinking about Lemma 2 is that it indicates that if one wants a “for all” guarantee on expansion, then one has to return to the regime of $m' = \mathcal{O}(k \log(n/k))$ measurements, as in “usual” compressive sensing.

Next, we formalize the “many $S(\mathbf{x})$ -leaf nodes” property defined in Property 4. Recall that for any set $S(\mathbf{x})$ of at most k nodes on the left of \mathcal{G} , we call any node on the right of \mathcal{G} an $S(\mathbf{x})$ -leaf node if it has exactly one neighbor in $S(\mathbf{x})$.

Lemma 3. *Let $S(\mathbf{x})$ be a set of k nodes on the left of \mathcal{G} such that the number of nodes neighbouring those in any $S'(\mathbf{x}) \subseteq S(\mathbf{x})$ is at least $2d/3$ times the size of $S'(\mathbf{x})$. Then at least a fraction $1/2$ of the nodes that are neighbours of any $S'(\mathbf{x}) \subseteq S(\mathbf{x})$ are $S'(\mathbf{x})$ -leaf nodes.*

Proof: Based on Lemma 1. Follows from a counting argument similar to those used in expander codes [33]. Proof in Appendix C.

C. Description of SHO-FA

Given a graph \mathcal{G} satisfying properties 1-4, we now describe our encoding and decoding procedure.

Measurement matrix:

Matrix structure and entries: The encoder’s *measurement matrix* A is chosen based on the structure of \mathcal{G} (recall that \mathcal{G} has n nodes on the left and m' nodes on the right). To begin with, the matrix A has $m = 2m'$ rows, and its non-zero values are unit-norm complex numbers. This choice of using complex numbers rather than real numbers in A is for notational convenience only. One can equally well choose a matrix A' with $m = 4m'$ rows, and replace each row of A with two consecutive rows in A' comprising respectively of the real and imaginary parts of rows of A . Since the components of \mathbf{x} are real numbers, hence there is a bijection between $A\mathbf{x}$ and $A'\mathbf{x}$ – indeed, consecutive pairs of elements in $A'\mathbf{x}$ are respectively the real and imaginary parts of the complex components of $A\mathbf{x}$. Also, as we shall see, the choice of unit-norm complex numbers ensures that “noise” due to finite precision arithmetic does not get “amplified”.

In particular, corresponding to node i on the right-hand side of \mathcal{G} , the matrix A has two rows. The j^{th} entries of the $(2i - 1)^{\text{th}}$ and $2i^{\text{th}}$ rows of A are respectively denoted $a_{i,j}^{(I)}$ and $a_{i,j}^{(V)}$ respectively. (The superscripts (I) and (V) respectively stand for *Identification* and *Verification*, for reasons that shall become clearer when we discuss the process to reconstruct \mathbf{x} .)

Identification entries: If \mathcal{G} has no edge connecting node j on the left with i on the right, then the *identification* entry $a_{i,j}^{(I)}$ is set to equal 0. Else, if there is indeed such an edge, $a_{i,j}^{(I)}$ is set to equal

$$a_{i,j}^{(I)} = e^{\iota j\pi/(2n)}. \quad (1)$$

(Here ι denotes the positive square root of -1 .) This entry $a_{i,j}^{(I)}$ can also be thought of as the weight of the edge in \mathcal{G} connecting j on the left with i on the right. In particular, the *phase* $j\pi/(2n)$ of $a_{i,j}^{(I)} = e^{\iota j\pi/(2n)}$ will be critical for our algorithm. As in Property 2 in Section II-B, our choice above guarantees distinct weights for all

edges connected to a node i on the right.

Verification entries: Whenever the identification entry $a_{i,j}^{(I)}$ equals 0, we choose to set the corresponding *verification entry* $a_{i,j}^{(V)}$ also to be zero. On the other hand, whenever $a_{i,j}^{(I)} \neq 0$, then we set $a_{i,j}^{(V)}$ to equal $e^{i\theta_{i,j}^{(V)}}$ for $\theta_{i,j}^{(V)}$ chosen uniformly at random from $[0, \pi/2]$ (with $\mathcal{O}(\log(k))$ bits of precision).¹⁵

Example 2: The matrix A corresponding to the graph \mathcal{G} in Example 1 is show in Figure 3.

$$A = \begin{bmatrix} e^{i\pi \cdot 0} & 0 & e^{i\pi/6} & e^{i\pi/3} & 0 \\ e^{i\theta_{1,1}} & 0 & e^{i\theta_{1,3}} & e^{i\theta_{1,4}} & 0 \\ \hline e^{i\pi \cdot 0} & e^{i\pi/6} & 0 & e^{i\pi/3} & e^{i\pi/2} \\ e^{i\theta_{2,1}} & e^{i\theta_{2,2}} & 0 & e^{i\theta_{2,4}} & e^{i\theta_{2,5}} \\ \hline e^{i\pi \cdot 0} & e^{i\pi/6} & e^{i\pi/3} & 0 & e^{i\pi/2} \\ e^{i\theta_{3,1}} & e^{i\theta_{3,2}} & e^{i\theta_{3,3}} & 0 & e^{i\theta_{3,5}} \\ \hline 0 & e^{i\pi \cdot 0} & e^{i\pi/6} & e^{i\pi/3} & e^{i\pi/2} \\ 0 & e^{i\theta_{4,2}} & e^{i\theta_{4,3}} & e^{i\theta_{4,4}} & e^{i\theta_{4,5}} \end{bmatrix}$$

Figure 3. This 8×5 matrix denotes the A corresponding to the graph \mathcal{G} . Note that its primary purpose is expository – clearly, 8 measurements (or indeed, 16 measurements over \mathbb{R}) to reconstruct a 2-sparse vector of length 5 is too many! Nonetheless, this is just an artifact of the fact that n in this example is small. In fact, according to our proofs, even as n scales to infinity, the number of measurements required to reconstruct a 2-sparse vector (or in general a k -sparse vector for constant k) remains constant! Also, note that we do not use the assignment for the identification entries $a_{i,j}^{(I)}$ specified in (1), since doing so would result in ugly and not very illuminating calculations in Example 3 below. However, as noted in Remark 1, this is not critical – it is sufficient that distinct entries in the identification rows of the matrix be distinct.

D. Reconstruction

Since the measurement matrix A has interspersed identification and verification rows, this induces corresponding interspersed *identification observations* $y_i^{(I)}$ and *verification observations* $y_i^{(V)}$ in the observation vector $\mathbf{y} = A\mathbf{x}$. Let $\mathbf{y}^{(I)} = \{y_i^{(I)}\}$ denote the length- m *identification vector* over \mathbb{C} , and $\mathbf{y}^{(V)} = \{y_i^{(V)}\}$ denote the length- m *verification vector* over \mathbb{C} .

Given the measurement matrix A and the observed $(\mathbf{y}^{(I)}, \mathbf{y}^{(V)})$ identification and verification vectors, the decoder’s task is to find *any* k -sparse vector $\hat{\mathbf{x}}$ such that $A\hat{\mathbf{x}}$ results in the corresponding identification and observation vectors. We shall argue below that if we succeed, then with high probability over A (specifically, over the verification entries of A), this $\hat{\mathbf{x}}$ must equal \mathbf{x} .

To find such a $\hat{\mathbf{x}}$ we design an iterative decoding scheme. This scheme starts by setting the initial guess for the reconstruction vector $\hat{\mathbf{x}}$ to the all-zero vector, and also identifies the set of non-zero indices of $\mathbf{y}^{(V)}$. It initializes the *neighborly set* as the set of non-zero *indices* of the *verification vector* $\mathbf{y}^{(V)}$, and initializes the *gap vector* as the *values* of the observation vector \mathbf{y} restricted to the neighbourly set. In the first iteration it then picks a uniformly random index i from the neighbourly set. Next, the decoder attempts to recover the signal value at some index $j \in \mathcal{S}(\mathbf{x})$ by looking at $y_i^{(I)}$ and “estimating” which j on the left of \mathcal{G} could have “caused the identification observation $y_i^{(I)}$ ”. If index i is not a $\mathcal{S}(\mathbf{x})$ -leaf node, the decoder does not succeed in reconstructing x_j , it declares the iteration as a failure, and starts the second iteration by again choosing a new uniformly random index i from the neighbourly set. On the other hand, if index i is a $\mathcal{S}(\mathbf{x})$ -leaf node, a signal value x_j will indeed be recovered (and “verified” using the verification entry $a_{i,j}^{(V)}$ and the verification observation $y_i^{(V)}$)¹⁶, then the algorithm will update the gap vectors by subtracting the “contribution” of the coordinate x_j to the measurements it influences (there are

¹⁵This choice of precision for the verification entries contributes one term to our expression for the precision of arithmetic required. As we argue later in Section II-K, this choice of precision guarantees that if a single identification step returns a value for x_j , this is indeed correct with probability $1 - o(1/k)$. Taking a union bound over $\mathcal{O}(k)$ indices corresponding to non-zero x_j gives us an overall $1 - o(1)$ probability of success.

¹⁶As Ronald W. Reagan liked to remind us, “*doveryai, no proveryai*”.

exactly three of them since the degree of the nodes on the left side of \mathcal{G} is 3), remove i from the neighborly set, and finally pick a new random index i from the neighbourly set for the next (second) iteration. The decoder performs the above operations repeatedly until $\hat{\mathbf{x}}$ has been completely recovered. We also show that (with high probability over A) in $\mathcal{O}(k)$ steps this process does indeed terminate.

Example 3: Figures 4–8 show a sample decoding process for the matrix A as in Example 2, and the observed vector \mathbf{y} shown in the figures. The example also demonstrates each of several possible scenarios the algorithm can find itself in, and how it deals with them.

E. Formal description of SHO-FA's reconstruction process

Our algorithm proceeds iteratively, and has $\mathcal{O}(k)$ overall (expected) number of iterations, with t being the variable indexing the iteration number.

- 1) We initialize by setting the *signal estimate vector* $\hat{\mathbf{x}}(1)$ to the all-zeros vector 0^n , and the *residual measurement identification/verification vectors* $\tilde{\mathbf{y}}^{(I)}(1)$ and $\tilde{\mathbf{y}}^{(V)}(1)$ to the decoder's observations $\mathbf{y}^{(I)}$ and $\mathbf{y}^{(V)}$. Let $\mathcal{D}(1)$, the initial neighborly set, be the set of indices i corresponding to non-zero locations of the initial verification vector $\mathbf{y}^{(V)}$, i.e., the set $\{i \leq m : \tilde{y}_i^{(V)}(1) \neq 0\}$. This step alone already takes $\mathcal{O}(k)$ steps, since merely reading \mathbf{y} to check for the zero locations of $\mathbf{y}^{(V)}$ takes that long.
- 2) The t^{th} decoding iteration accepts as its input the t^{th} signal estimate vector $\hat{\mathbf{x}}(t)$, the t^{th} neighbourly set $\mathcal{D}(t)$, and the t^{th} residual measurement identification/verification vectors $(\tilde{\mathbf{y}}^{(I)}(t), \tilde{\mathbf{y}}^{(V)}(t))$. In $\mathcal{O}(1)$ steps it outputs the $(t+1)^{\text{th}}$ signal estimate vector $\hat{\mathbf{x}}(t+1)$, the $(t+1)^{\text{th}}$ neighbourly set $\mathcal{D}(t+1)$, and the $(t+1)^{\text{th}}$ residual measurement identification/verification vectors $(\tilde{\mathbf{y}}^{(I)}(t+1), \tilde{\mathbf{y}}^{(V)}(t+1))$ after the performing the following steps sequentially (each of which takes at most a constant number of atomic steps):
- 3) *Pick a random $i(t)$:* The decoder picks an element $i(t)$ uniformly at random from the t^{th} neighborly set $\mathcal{D}(t)$.
- 4) *Compute angles $\theta^{(I)}(t)$ and $\theta^{(V)}(t)$:* Let the *current identification and verification angles* be defined respectively as the phases of the residual identification and verification entries being considered in that step, as follows:

$$\begin{aligned}\theta^{(I)}(t) &\triangleq \angle \left(\tilde{y}_{i(t)}^{(I)}(t) \right), \\ \theta^{(V)}(t) &\triangleq \angle \left(\tilde{y}_{i(t)}^{(V)}(t) \right).\end{aligned}$$

Here $\angle(\cdot)$ computes the phase of a complex number (up to $\mathcal{O}(\max\{\log n/k, \log(k)\})$ bits of precision)¹⁷.

- 5) *Check if the current identification and verification angles correspond to a valid and unique x_j :* For this, we check at most two things (both calculations are done up to the precision specified in the previous step).
 - a) First, we check if $j(t) \triangleq \theta^{(I)}(t)(2n/\pi)$ is an integer, and the corresponding j^{th} element of the i^{th} row is non-zero. If so, we have “tentatively identified” that the i^{th} component of $\tilde{\mathbf{y}}$ is a leaf-node of the currently unidentified non-zero components of \mathbf{x} , and in particular is connected to the $j(t)^{\text{th}}$ node on the left, and the algorithm proceeds to the next step below. If not, we simply increment t by 1 and return to Step (3).
 - b) Next, we verify our estimate from the previous step. If $a_{i(t),j(t)}^{(V)} \tilde{y}_{i(t)}^{(I)}(t) / a_{i(t),j(t)}^{(I)} = \tilde{y}_{i(t)}^{(V)}(t)$, the verification test passes, and the algorithm proceeds to the next step below. If not, we simply increment t by 1 and return to Step (3).
- 6) *Update $\hat{\mathbf{x}}(t+1)$, $\mathcal{D}(t+1)$, $\tilde{\mathbf{y}}^{(I)}(t+1)$, and $\tilde{\mathbf{y}}^{(V)}(t+1)$:* In particular, at most 3 components of each of these vectors need to be updated. Specifically, $\hat{x}_{j(t)}(t+1)$ equals $\tilde{y}_{i(t)}^{(I)}(t) / a_{i(t),j(t)}^{(I)}$. The (at most three) neighbours of $\hat{x}_{j(t)}(t)$ are removed from the neighbourly set $\mathcal{D}(t)$ to get the neighbourly set $\mathcal{D}(t+1)$. And finally (three) values each of $\tilde{\mathbf{y}}^{(I)}(t+1)$ and $\tilde{\mathbf{y}}^{(V)}(t+1)$ are updated from those of $\tilde{\mathbf{y}}^{(I)}(t)$ and $\tilde{\mathbf{y}}^{(V)}(t)$ (those corresponding to the neighbours of $\hat{x}_{j(t)}(t)$) by subtracting out $\hat{x}_{j(t)}(t)$ multiplied by the appropriate coefficients of A .
- 7) *Termination:* The algorithm stops when the neighbourly set is empty, and outputs the last $\hat{\mathbf{x}}(t)$.

¹⁷Roughly, the former term guarantees that the identification angle is calculated precisely enough, and the latter that the verification angle is calculated precisely enough.




Figure 4. *Initialization*: The (true) \mathbf{x} equals $(0, 1, 0, 1, 0)$ (and hence $\mathcal{S}(\mathbf{x}) = \{2, 4\}$). Also note that nodes 1 and 3 on the right of \mathcal{G} are $\mathcal{S}(\mathbf{x})$ -leaf nodes, as defined in Property 4. However, all of this is unknown to the decoder *a priori*. The decoder sets the (starting) estimate $\hat{\mathbf{x}}(0)$ of the reconstruction vector $\hat{\mathbf{x}}$ to the all-zeros vector. The (starting) gap vector $\tilde{\mathbf{y}}$ is set to equal \mathbf{y} , which in turn equals the corresponding 4 pairs of identification and verification observations on the right-hand side of \mathcal{G} . The specific values of $\theta_{i,j}^{(V)}$ in the verification observations do not matter currently – all that matters is that given \mathbf{x} , each of the four verification observations are non-zero (with high probability over the choices of $\theta_{i,j}^{(V)}$). Hence the (starting) value of the *neighbourly* set equals $\{1, 2, 3, 4\}$. This step takes $\mathcal{O}(k)$ number of steps, just to initialize the neighbourly set. By the end of the decoding algorithm (if it runs successfully), the tables will be turned – all the entries on the right of \mathcal{G} will equal zero, and (at most) k entries on the left of \mathcal{G} will be non-zero.

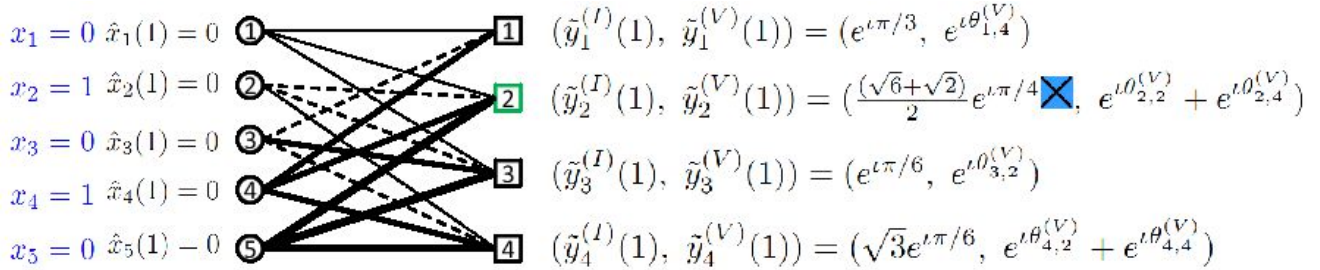


Figure 5. *Failed identification*: In this, the first iteration, the decoder randomly picks the index $i = 2$ from the neighbourly set $\{1, 2, 3, 4\}$, and checks the phase of the corresponding gap vector identification observation $\tilde{y}_2^{(I)}$. Since this equals $\pi/4$, which is *not* in the set of possible phases in the 2^{nd} identification row of A (which are all multiples of $\pi/6$), the decoder declares a failure in this iteration. In particular, the decoder is unable to (currently) use $\tilde{y}_2^{(I)}$ to identify a non-zero location of \mathbf{x} , since the second node on the right of \mathcal{G} is not a $\mathcal{S}(\mathbf{x})$ -leaf node. So far, the reconstruction vector $\hat{\mathbf{x}}$, the gap vector $\tilde{\mathbf{y}}$, and the neighbourly set $\{1, 2, 3, 4\}$ are all still unchanged. This entire iteration takes a constant number of steps.




Figure 6. *Passed identification, failed verification*: In the second iteration, a potentially more serious failure could happen. In particular, suppose the decoder randomly picks the index $i = 4$ from the neighbourly set $\{1, 2, 3, 4\}$ (note that 4 is *also* not a $\mathcal{S}(\mathbf{x})$ -leaf node), and checks the phase of the corresponding gap vector identification observation $\tilde{y}_4^{(I)}$, it just so happens that the value of \mathbf{x} is such that this corresponds to a phase of $\pi/6$. But as can be seen from the matrix in Figure 3, for $i = 4$ this corresponds to $a_{i,j}^{(I)}$ for $j = 3$. Hence the decoder would make a “false identification” of $j = 3$, and estimate that \hat{x}_3 equals the magnitude of $\tilde{y}_4^{(I)}$, which would equal $\sqrt{3}$. This is where the verification entries and verification observations save the day. Recall that each verification entry is chosen uniformly at random (with sufficient bit precision) from $[0, 2\pi)$, independently of both \mathbf{x} and the other entries of A . Hence the probability that $\sqrt{3}$ (the misdirected value of \hat{x}_3) times the corresponding verification entry $a_{4,3}^{(V)}$ equals $\tilde{y}_4^{(V)}$ is “small”. Hence the decoder in this case too declares the iteration to be a failure, and leaves the reconstruction vector $\hat{\mathbf{x}}$, the gap vector $\tilde{\mathbf{y}}$, and the neighbourly set $\{1, 2, 3, 4\}$ all still unchanged. This entire iteration takes a constant number of steps.




Figure 7. *Passed identification, passed verification*: Now, in the third iteration, suppose the decoder randomly picks the index $i = 1$ from the neighbourly set $\{1, 2, 3, 4\}$ (note that 1 is a $\mathcal{S}(\mathbf{x})$ -leaf node). In this case, the phase of the corresponding gap vector identification observation $\tilde{y}_1^{(I)}$ equals $\pi/3$. As can be seen from the matrix in Figure 3, for $i = 1$ this corresponds to $a_{i,j}^{(I)}$ for $j = 4$. Hence the decoder makes a “correct identification” of $j = 4$, and estimates (also correctly) that \hat{x}_4 equals the magnitude of $\tilde{y}_1^{(I)}$, which equals 1. On checking with the verification entry, the decoder observes also that 1 (the detected value of \hat{x}_4) times the corresponding verification entry $a_{1,4}^{(V)}$ equals $\tilde{y}_1^{(V)}$. Hence it updates the value of \hat{x}_4 to 1, the neighbourly set to $\{2, 3, 4\}$, and $\tilde{\mathbf{y}}$ to the values shown (only the three indices 1, 3 and 4 on the right need to be changed). At this point, note that $\mathcal{S}'(\mathbf{x})$ also changes from $\{2, 4\}$ to the singleton set $\{4\}$. This entire iteration takes a constant number of steps.




Figure 8. *Termination*: In the fourth iteration, the decoder randomly picks $i = 4$ from the neighbourly set $\{2, 3, 4\}$. Recall that in the second iteration this choice of i did not aid in decoding. However, now that node 4 on the right of \mathcal{G} has been “cleaned up”, it is now a leaf node for $\mathcal{S}'(\mathbf{x})$. This demonstrates the importance of not “throwing away” information which seems useless at some point in time. Hence, analogously to the process in Figure 7, the decoder estimates the value of \hat{x}_2 to 1, updates the neighbourly set to the empty set, and $\tilde{\mathbf{y}}$ to the all-zero vector (all in a constant number of steps). Since the gap vector is zero, this indicates to the decoder that it should output $\hat{\mathbf{x}}$ as its estimate of \mathbf{x} , and terminate.

F. Expected number of iterations

We first argue that, with a constant probability, each iteration result in recovering a new non-zero coordinate from \mathbf{x} . Towards this, for each $t = 1, 2, \dots$, let $\mathcal{S}(t)$ be the support of $\mathbf{x} - \hat{\mathbf{x}}(t)$. Note that $\mathcal{D}(t) = N(\mathcal{S}(t))$ and $\mathcal{S}(\mathbf{x}) = \mathcal{S}(1) \supseteq \mathcal{S}(2) \supseteq \dots$

Then, according to Lemma 3 and the way we generate the measurement matrix A , with a high probability, for each t , the probability that there exists a node $i(t)$ in $\tilde{\mathbf{y}}^{(I)}(t)$ so that it is an $\mathcal{S}(t)$ -leaf node is lower bounded by $1/2$. Consequently, exactly one non-zero coordinate in $\mathcal{S}(t)$ completely determines $\tilde{\mathbf{y}}_{i(t)}^{(I)}(t)$ and $\tilde{\mathbf{y}}_{i(t)}^{(V)}(t)$. The algorithm identifies this coordinate as j for the t^{th} iteration and at the end of iteration, recovers \hat{x}_j . Thus, whenever $i(t)$ is an $\mathcal{S}(t)$ -leaf node, the set of recovered coordinates increases by 1. When $i(t)$ is not an $\mathcal{S}(t)$ -leaf node, our reconstruction process wastes one iteration and will start another iteration by picking another node from the neighborly set $\mathcal{D}(t)$ uniformly at random. Hence, the operations among different iterations are independent, and each iteration succeeds with probability $1/2$.

Since there are at most k non-zero coordinates in \mathbf{x} , the number of iterations before the algorithm terminates follows a Pascal distribution with parameters $(k, 1/2)$. The expected number of iterations is then simply $2k$.

G. Correctness

Next, we show that $\hat{\mathbf{x}} = \mathbf{x}$ with a high probability. To show this, it suffices to show that each non-zero update to the estimate $\hat{\mathbf{x}}(t)$ sets a previously untouched coordinate to the correct value with a high probability.

Note that if $i(t)$ is a leaf node for $\mathcal{S}(t)$, and if all non-zero coordinates of $\hat{\mathbf{x}}(t)$ are equal to the corresponding coordinates in \mathbf{x} , then the decoder correctly identifies the parent node $j(t) \in \mathcal{S}(t)$ for the leaf node $i(t)$ as the unique coordinate that passes the phase identification and verification checks.

Thus, the t^{th} iteration ends with an erroneous update only if

$$\angle\left(\sum_{p \in N(\{i(t)\})} x_p e^{j\theta_{i(t),p}^{(I)}}\right) = \theta_{i(t),j(t)}^{(I)}$$

for some j such that there are more than one non-zero terms in the summation on the left.

$$\angle\left(\sum_{p \in N(\{i(t)\})} x_p e^{j\theta_{i(t),p}^{(V)}}\right) = \theta_{i(t),j(t)}^{(V)}$$

Since $V(i(t), j)$ is drawn uniformly at random from $\{1, 2, \dots, \lceil 4n \rceil\}$, the probability that the second equality holds with more than one non-zero term in the summation on the left is at most $1/(4n)$. The above analysis gives an upper bound on the probability of incorrect update for a single iteration to be $1/(4n)$. Finally, as the total number of updates is at most k , by applying a union bound over the updates, the probability of incorrect decoding is upper bounded by $k/4n$. Since $k = o(n)$ by assumption, it follows that the error probability vanishes as n and k grow without bound.

H. Remarks on the Reconstruction process for exactly k -sparse signals

We elaborate on these choices of entries of A in the remarks below, which also give intuition about the reconstruction process outlined in Section II-E.

Remark 1: In fact, it is not critical that (1) be used to assign the identification entries. As long as j can be "quickly" (computationally efficiently) identified from the phases of $a_{i,j}^{(I)}$ (as outlined in Remark 2 below, and specified in more detail in Section II-E), this suffices for our purpose. This is the primary reason we call these entries identification entries.

Remark 2: The reason for the choice of phases specified in (1) is as follows. Suppose $\mathcal{S}(\mathbf{x})$ corresponds to the support (set of non-zero values) of \mathbf{x} . Suppose y_i corresponds to a $\mathcal{S}(\mathbf{x})$ -leaf node, then by definition $y_i^{(I)}$ equals $a_{i,j}^{(I)} x_j$ for some j in $\{1, \dots, n\}$ (if y_i corresponds to a $\mathcal{S}(\mathbf{x})$ -non-leaf node, then in general $y_i^{(I)}$ depends on two or more x_j). But x_j is a real number. Hence examining the phase of y_i enables one to efficiently compute $j\pi/(2n)$, and hence j . It also allows one to recover the magnitude of x_j , simply by computing the magnitude of y_i .

Remark 3: The choice of phases specified in (1) divides the set of allowed phases (the interval $[0, \pi/2]$) into n distinct values. Two things are worth noting about this choice.

- 1) We consider the interval $[0, \pi/2]$ rather than the full range $[0, 2\pi)$ of possible phases since we wish to use the phase measurements to *also* recover the sign of x_j s. If the phase of y_i falls within the interval $[0, \pi/2]$, then (still assuming that y_i corresponds to a $\mathcal{S}(\mathbf{x})$ -leaf node) x_j must have been positive. On the other hand, if the phase of y_i falls within the interval $[\pi, 3\pi/2]$, then x_j must have been negative. (It can be directly verified that the phase of a $\mathcal{S}(\mathbf{x})$ -leaf node y_i can never be outside these two intervals – this wastes roughly half of the set of possible phases we could have used for identification purposes, but it makes notation easier.
- 2) The choice in (1) divides the interval $[0, \pi/2]$ into n distinct values. However, in expectation over \mathcal{G} the actual number of non-zero entries in a row of A is $\mathcal{O}(n/k)$, so on average one only needs to choose $\mathcal{O}(n/k)$ distinct phases in (1), rather than the worst case n number of values. This has the advantage that one only needs $\mathcal{O}(\log(n/k))$ bits of precision to specify distinct phase values (and in fact we claim that this is the level of precision required by our algorithm). However, since we analyze only left-regular \mathcal{G} , the degrees of nodes on the right will in general vary stochastically around this expected value. If k is “somewhat large” (for instance $k = \Omega(n)$), then the degrees will not be very tightly concentrated around their mean. One way around this is to choose \mathcal{G} uniformly at random from the set of bipartite graphs with n nodes (each of degree d) on the left and m nodes (each of degree dn/m) on the right. This would require a more intricate proof of the $\mathcal{S}'(\mathbf{x})$ -expansion property defined in Property 3 and proved in Lemma 1. For the sake of brevity, we omit this proof here.

Remark 4: In fact, the recent work of [32] demonstrates an alternative analytical technique (bypassing the expansion arguments outlined in this work), involving analysis of properties of the “2-core” of random hyper-graphs, that allows for a tight characterization of the number of measurements required by SHO-FA to reconstruct \mathbf{x} from \mathbf{y} and A , rather than the somewhat loose (though order-optimal) bounds presented in this work. Since our focus in this work is a simple proof of order-optimality (rather than the somewhat more intricate analysis required for the tight characterization) we again omit this proof here.¹⁸

I. Database query

A useful property of our construction of the matrix A is that any desired signal component x_j can be reconstructed with a constant probability given the measurement vector $\mathbf{y} = A\mathbf{x}$ in a constant time. The following Lemma makes this precise. The proof follows from a simple probabilistic argument and is included in Appendix D.

Lemma 4. *Let \mathbf{x} be k -sparse. Let $j \in \{1, 2, \dots, n\}$ and let $A \in \mathbb{C}^{ck \times n}$ be randomly drawn according to SHO-FA. Then, there exists an algorithm \mathcal{A} such that given inputs (j, \mathbf{y}) , \mathcal{A} produces an output \hat{x}_j with probability at least $(1 - (d/c)^d)$ such that $\hat{x}_j = x_j$ with probability $(1 - o(1/k))$.*

J. SHO-FA for sparse vectors in different bases

In the setting of SHO-FA we consider k -sparse input vectors \mathbf{x} . In fact, we also can deal with the case that \mathbf{x} is sparse in a certain basis that is known *a priori* to the decoder¹⁹, say Ψ , which means that $\mathbf{x} = \Psi\mathbf{w}$ where \mathbf{w} is a k -sparse vector. Specifically, in this case we write the measurement vector as $\mathbf{y} = B\mathbf{x}$, where $B = A\Psi^{-1}$. Then, $\mathbf{y} = A\Psi^{-1}\Psi\mathbf{w} = A\mathbf{w}$, where A is chosen on the structure of the \mathcal{G} and \mathbf{w} is a k -sparse vector. We can then apply SHO-FA to reconstruct \mathbf{w} and consequently $\mathbf{x} = \Psi\mathbf{w}$. What has been discussed here covers the case where \mathbf{x} is sparse itself, for which we can simply take $\Psi = I$ and $\mathbf{x} = \mathbf{w}$.

K. Information-theoretically optimal number of bits

We recall that the reconstruction goal for SHO-FA is to reconstruct \mathbf{x} up to relative error 2^{-P} . That is,

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_1 / \|\mathbf{x}\|_1 \leq 2^{-P}.$$

We first present a sketch of an information-theoretic lower bound of $\Omega(k(P + \log n))$ bits holds for any algorithm that outputs a k -sparse vector that achieves this goal with high probability.

To see this is true, consider the case where the locations of k non-zero entries in \mathbf{x} are chosen uniformly at random among all the n , entries and the value of each non-zero entry is chosen uniformly at random from the set $\{1, \dots, 2^P\}$. Then recovering even the support requires at least $\log(2^{kP} \binom{n}{k})$ bits, which is $\Omega(k \log(n/k))$.²⁰

¹⁸We thank the anonymous reviewers who examined a previous version of this work for pointing out the extremely relevant techniques of [32] and [34] (though the problems considered in those works were somewhat different).

¹⁹For example, “smooth” signals are sparse in the Fourier basis and “piecewise smooth” signals are sparse in wavelet bases.

²⁰Stirling’s approximation(c.f. [47, Chapter 1]) is used in bounding from below the combinatorial term $\binom{n}{k}$.




Figure 9. *An example of a physical system that “naturally” generates ensembles of sparse A that SHO-FA can use:* Suppose there are k cellphones (out of a set of n possible different cellphones in the whole world) in a certain neighbourhood that has a base-station. The goal is for the j -th cellphone to communicate its information (x_j) to the base-station at least once per *frame* of ck consecutive time-slots. The challenge is to do so in a distributed manner, since multiple cellphones transmitting at the same time i would result in a linear combination $y_i = \sum_j a_{ij}x_j$ of their transmissions reaching the base-station, where a_{ij} corresponds to the channel gain from the cellphone j to the base-station during time-slot i . Each cellphone transmits x_j to the base-station a constant (d) number of times in each frame – the set of d time-slots in each frame that cellphone j transmits in is chosen by j uniformly at random from the set of all $\binom{ck}{d}$ sets of slots.

Also, at least a constant fraction of the k non-zero entries of \mathbf{x} must be correctly estimated to guarantee the desired relative error. Hence $\Omega(k(P + \log n))$ is a lower bound on the measurement bit-complexity.

The following arguments show that the total number of bits used in our algorithm is information-theoretically order-optimal for any $k = \mathcal{O}(n^{1-\Delta})$ (for any $\Delta > 0$). First, to represent each non-zero entry of \mathbf{x} , we need to use arithmetic of $\Omega(P + \log(k))$ bit precision. Here the P term is so as to attain the required relative error of reconstruction, and the $\log(k)$ term is to take into account the error induced by finite-precision arithmetic (say, for instance, by floating point numbers) in $\mathcal{O}(k)$ iterations (each involving a constant number of finite-precision additions and unit-magnitude multiplications). Second, for each identification step, we need to use $\Omega(\log(n) + \log(k))$ bit-precision arithmetic. Here the $\log(n)$ term is so that the identification measurements can uniquely specify the locations of non-zero entries of \mathbf{x} . The $\log(k)$ term is again to take into account the error induced in $\mathcal{O}(k)$ iterations. Third, for each verification step, the number of bits we use are $3\log(k)$. Here, by the Schwartz-Zippel Lemma [48], [49], $2\log(k)$ bit-precision arithmetic guarantees that each verification step is valid with probability at least $1 - 1/k^2$ – a union bound over all $\mathcal{O}(k)$ verification steps guarantees that all verification steps are correct with probability at least $1 - \mathcal{O}(1/k)$. Therefore, the total number of bits needed by SHO-FA $\mathcal{O}(k(\log(n) + P))$. As claimed, this matches, up to a constant factor, the lower bound sketched above.

L. Universality

While the ensemble of matrices $\{A\}$ we present above has carefully chosen identification entries, and all the non-zero verification entries have unit magnitude, we argue that in fact the implicit ideas underlying SHO-FA work for significantly more general ensembles of matrices. In particular, Property 1 only requires that the graph \mathcal{G} underlying A be “sparse”, with a constant number of non-zero entries per column. Property 1 only requires that each non-zero entry in each row be distinct – which is guaranteed with high probability, for instance, if each entry is chosen *i.i.d* from any distribution with sufficiently large support. An example of such a scenario is shown in Figure 9. This naturally motivates the application of SHO-FA to a variety of scenarios, for e.g., neighbor discovery in wireless communication [38].

III. APPROXIMATE RECONSTRUCTION IN THE PRESENCE OF NOISE

A prominent aspect of the design presented in the previous section is that it relies on exact determination of all the phases as well as magnitudes of the measurement vector $A\mathbf{x}$. In practice, however, we often desire that the measurement and reconstruction be robust to corruption both before and during measurements. In this section, we show that our design may be modified slightly such that with a suitable decoding procedure, the reconstruction is robust to such "noise".

We consider the following setup. Let $\mathbf{x} \in \mathbb{R}^n$ be a k -sparse signal with support $\mathcal{S}(\mathbf{x}) = \{j : x_j \neq 0\}$. Let $\mathbf{z} \in \mathbb{R}^n$ have support $\{1, 2, \dots, n\} \setminus \mathcal{S}(\mathbf{x})$ with each z_j distributed according to a Gaussian distribution with mean 0 and variance σ_z^2 . Denote the measurement matrix by $A \in \mathbb{C}^{m \times n}$ and the measurement vector by $\mathbf{y} \in \mathbb{C}^m$. Let $\mathbf{e} \in \mathbb{C}^m$ be the measurement noise with e_i distributed as a Complex Gaussian with mean 0 and variance σ_e^2 along each axis. \mathbf{y} is related to the signal as

$$\mathbf{y} = A(\mathbf{x} + \mathbf{z}) + \mathbf{e}.$$

We propose a design procedure for A satisfying the following properties.

Theorem 2. *Let $k = \mathcal{O}(n^{1-\Delta})$ for some $\Delta > 0$. There exists a reconstruction algorithm SHO-FA for $A \in \mathbb{C}^{m \times n}$ such that*

- (i) $m = ck$
- (ii) SHO-FA consists of at most $4k$ iterations, each involving a constant number of arithmetic operations with a precision of $\mathcal{O}(\log n)$ bits.
- (iii) With probability $1 - o(1/k)$ over the design of A and randomness in \mathbf{e} and \mathbf{z} ,

$$\|\hat{\mathbf{x}} - \mathbf{x}\|_1 \leq C (\|\mathbf{z}\|_1 + (\log k)^2 \|\mathbf{e}\|_1)$$

for some $C = C(\sigma_z, \sigma_e) > 0$.

Recall that in the exactly k -sparse case, the decoding in t -th iteration relies on first finding an $\mathcal{S}(t)$ -leaf node, then decoding the corresponding signal coordinate and updating the undecoded measurements. In this procedure, it is critical that each iteration operates with low reconstruction errors as an error in an earlier iteration can propagate and cause potentially catastrophic errors. In general, one of the following events may result in any iteration ending with a decoded signal value that is far from the true signal value:

- (a) The decoder picks an index outside the set $\{i : (A\mathbf{x})_i \neq 0\}$, but in the set $\{i : (A(\mathbf{x} + \mathbf{z}) + \mathbf{e})_i \neq 0\}$.
- (b) The decoder picks an index within the set $\{i : (A\mathbf{x})_i \neq 0\}$ that is also a leaf for \mathcal{S} with parent node j , but the presence of noise results in the decoder identifying (and verifying) a node $j' \neq j$ as the parent and, subsequently, incorrectly decoding the signal at j' .
- (c) The decoder picks an index within the set $\{i : (A\mathbf{x})_i \neq 0\}$ that is not a leaf for \mathcal{S} , but the presence of noise results in the decoder identifying (and verifying) a node j as the parent and, subsequently, incorrectly decoding the signal at j .
- (d) The decoder picks an index within the set $\{i : (A\mathbf{x})_i \neq 0\}$ that is a leaf for \mathcal{S} with parent node j , which it also identifies (and verifies) correctly, but the presence of noise introduces a small error in decoding the signal value. This error may also propagate to the next iteration and act as "noise" for the next iteration.

To overcome these hurdles, our design takes the noise statistics into account to ensure that each iteration is resilient to noise with a high probability. This achieved through several new ideas that are presented in the following for ease of exposition. Next, we perform a careful analysis of the corresponding decoding algorithm and show that under certain regularity conditions, the overall failure probability can be made arbitrarily small to output a reconstruction that is robust to noise. Key to this analysis is bounding the effect of propagation of estimation error as the decoder steps through the iterations.²¹

²¹For simplicity, the analysis presented here relies only on an upper bound on the length of the path through which the estimation error introduced in any iteration can propagate. This bound follows from known results on size of largest components in sparse hypergraphs [50]. We note, however, that a tighter analysis that relies on a finer characterization of the interaction between the size of these components and the contribution to total estimation error may lead to better bounds on the overall estimation error. Indeed, as shown in [34], such an analysis enables us to achieve a tighter reconstruction guarantee of the form $\|\mathbf{x} - \hat{\mathbf{x}}\|_1 = \mathcal{O}(\|\mathbf{z}\|_1 + \|\mathbf{e}\|_1)$

A. Key ideas

1) *Truncated reconstruction:* We observe that in the presence of noise, it is unlikely that signal values whose magnitudes are comparable to that of the noise values can be successfully recovered. Thus, it is futile for the decoder to try to reconstruct these values as long as the overall penalty in l_1 -norm is not high. The following argument shows that this is indeed the case. Let

$$\mathcal{S}_\delta(\mathbf{x}) = \{j : |x_j| < \delta/k\}. \quad (2)$$

and let $\mathbf{x}_{\mathcal{S}_\delta}$ be the vector defined as

$$(x_{\mathcal{S}_\delta})_j = \begin{cases} 0 & j \notin \mathcal{S}_\delta(\mathbf{x}) \\ x_j & j \in \mathcal{S}_\delta(\mathbf{x}). \end{cases}$$

Similarly, define $\mathbf{x}_{\mathcal{S}_\delta^c}$ which has non-zero entries only within the set $\mathcal{S}(\mathbf{x}) \setminus \mathcal{S}_\delta(\mathbf{x})$. The following sequence of inequalities shows that the total l_1 norm of $\mathbf{x}_{\mathcal{S}_\delta}$ is small:

$$\begin{aligned} \|\mathbf{x}_{\mathcal{S}_\delta}\|_1 &= \sum_{j \in \mathcal{S}_\delta(\mathbf{x})} |x_j| \\ &\leq |\mathcal{S}_\delta(\mathbf{x})| \frac{\delta}{k} \\ &\leq |\mathcal{S}(\mathbf{x})| \frac{\delta}{k} \\ &= \delta. \end{aligned} \quad (3)$$

Further, as an application of triangle inequality and the bound in (3), it follows that

$$\begin{aligned} \|\hat{\mathbf{x}} - \mathbf{x}\|_1 &= \|\hat{\mathbf{x}} - \mathbf{x}_{\mathcal{S}_\delta^c} - \mathbf{x}_{\mathcal{S}_\delta}\|_1 \\ &\leq \|\hat{\mathbf{x}} - \mathbf{x}_{\mathcal{S}_\delta^c}\|_1 + \|\mathbf{x}_{\mathcal{S}_\delta}\|_1 \\ &\leq \|\hat{\mathbf{x}} - \mathbf{x}_{\mathcal{S}_\delta^c}\|_1 + \delta \end{aligned} \quad (4)$$

Keeping the above in mind, we rephrase our reconstruction objective to satisfy the following criterion with a high probability:

$$\|\hat{\mathbf{x}} - \mathbf{x}_{\mathcal{S}_\delta^c}\|_1 \leq C_1(\|\mathbf{z}\|_1 + k^2\|\mathbf{e}\|_1), \quad (5)$$

while simultaneously ensuring that our choice of parameter δ satisfies

$$\delta < C_2\|\mathbf{z}\|_1 \quad (6)$$

for some C_2 , with a high probability.

2) *Phase quantization:* In the noisy setting, even when i is a leaf node for $\mathcal{S}(\mathbf{x})$, the phase of y_i may differ from the phase assigned by the measurement. This is geometrically shown in Figure 11a for a measurement matrix A' . To overcome this, we modify our decoding algorithm to work with "quantized" phases, rather than the actual received phases. The idea behind this is that if i is a leaf node for $\mathcal{S}(\mathbf{x})$, then quantizing the phase to one of the values allowed by the measurement identifies the correct phase with a high probability. The following lemma facilitates this simplification.

Lemma 5 (Almost bounded phase noise). *Let $\mathbf{x}, \mathbf{z} \in \mathbb{R}^n$ with $|x_j| > \delta/k$ for each j . Let $A' \in \mathbb{C}^{m' \times n}$ be a complex valued measurement matrix with the underlying graph \mathcal{G} . Let i be a leaf node for $\mathcal{S}(\mathbf{x})$. Let $\Delta\theta_i = |\angle y_i - \angle(A'\mathbf{x})_i|$. Then, for every $\alpha > 0$,*

$$E_{\mathbf{z}, \mathbf{e}}(\Delta\theta_i) \leq \sqrt{\frac{2\pi k^2(dn\sigma_z^2/c k + \sigma_e^2)}{\delta^2}}$$

and

$$\Pr_{\mathbf{z}, \mathbf{e}}\left(\Delta\theta_i > \alpha E_{\mathbf{z}, \mathbf{e}}(\Delta\theta_i)\right) < \frac{1}{2}e^{-(\alpha^2/\pi)}.$$

Proof: See Appendix E. ■

For a desired error probability ϵ' , the above lemma stipulates that it suffices to let $\alpha = (1/2)\log(1/2\epsilon')$. We examine the effect of phase noise in more detail in Appendix F.




Figure 10. The black curve corresponds to the magnitudes of $\mathbf{x} + \mathbf{z}$ (for ease of visual presentation, the components of \mathbf{x} have been sorted in decreasing order of magnitude and placed in the first k components of the signal, but the components of \mathbf{z} are unsorted. The blue curve corresponds to our reconstruction $\hat{\mathbf{x}}$ of \mathbf{x} . Note that we only attempt to reconstruct components of \mathbf{x} that are “sufficiently large” (that is, we make no guarantees about correct reconstruction of components of \mathbf{x} in $S_\delta(\mathbf{x})$, i.e., those components of \mathbf{x} that are smaller than some “threshold” δ/k . Here δ is a parameter of code-design to be specified later. As shown in Section III-A1, as long as δ is not “too large”, this relaxation does not violate our relaxed reconstruction criteria (5).

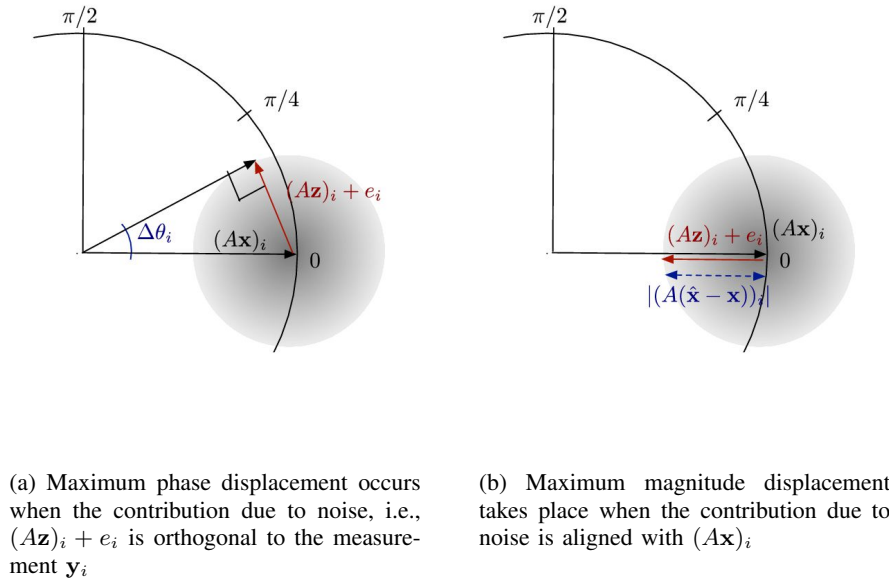


Figure 11. The effect of noise on a measurement output

3) *Repeated measurements*: Our algorithm works by performing a series of $\Gamma \geq 1$ identification and verification measurements in each iteration instead of a single measurement of each type as done in the exactly k -sparse case. The idea behind this is that, in the presence of noise, even though a single set of identification and verification measurements cannot exactly identify the coordinate j from the observed y_i , it helps us narrow down the set of coordinates j that can possibly contribute to give the observed phase. Performing measurements repeatedly, each time with a different measurement matrix, helps us identify a single j with a high probability.

We implement the above idea by first mapping each $j \in \{1, 2, \dots, n\}$ to its Γ -digit representation in base $G = \{0, 1, \dots, \lceil n^{1/\Gamma} \rceil - 1\}$. For each $j \in \{1, 2, \dots, n\}$, let $g(j) = (g_1(j), g_2(j), \dots, g_\Gamma(j))$ be the Γ -digit representation of j . Next, perform one pair of identification and verification measurements (and corresponding phase reconstructions), each of which is intended to distinguish exactly one of the digits. In our construction, we only need a constant number of such phase measurements per iteration. See Fig 12 for an illustrating example.




Figure 12. If we were to distinguish each j from 1 to 8 by a different phase, the decoder can tolerate at most $\pi/14$ phase displacement for any output y_i . Instead, we first represent each $j = 1, 2, \dots, 8$ by a three-length binary vector. Next, we perform three sets of phase assignments – one for each digit. It is easily seen that by allowing multiple measurements, the noise tolerance for the decoder increases.

B. Description of measurements

As in the exactly k -sparse case, we start with a randomly drawn left regular bipartite graph \mathcal{G} with n nodes on the left and m' nodes on the right.

Measurement matrix: The measurement matrix $A \in \mathbb{C}^{2m'\Gamma \times n}$ is chosen based on the graph \mathcal{G} . The rows of A are partitioned into m' groups, with each group consisting of 2Γ consecutive rows. The j -th entries of the rows $2(i-1)\Gamma+1, (i-1)\Gamma+2, \dots, 2i\Gamma$ are denoted by $a_{i,j}^{(I,1)}, a_{i,j}^{(I,2)}, \dots, a_{i,j}^{(I,\Gamma)}, a_{i,j}^{(V,1)}, a_{i,j}^{(V,2)}, \dots, a_{i,j}^{(V,\Gamma)}$ respectively. In the above notation, I and V are used to refer to identification and verification measurements.

For ease of notation, for each $\gamma = 1, 2, \dots, \Gamma$, we use $A^{(I,\gamma)}$ (resp. $A^{(V,\gamma)}$) to denote the sub-matrix of A whose (i,j) -th entry is $a_{i,j}^{(I,\gamma)}$ (resp. $a_{i,j}^{(V,\gamma)}$).

We define the γ -th identification matrix $A^{(I,\gamma)}$ as follows. For each (i,j) , if the graph \mathcal{G} does not have an edge connecting i on the right to j on the left, then $a_{i,j}^{(I,\gamma)} = 0$. Otherwise, we set $a_{i,j}^{(I,\gamma)}$ to be the unit-norm complex number

$$a_{i,j}^{(I,\gamma)} = e^{i g_\gamma(j) \pi / 2 (|\mathbf{G}| - 1)}.$$

Remark: Note here that the construction for the exactly k -sparse case can be recovered by setting $\Gamma = 1$, which results in $\mathbf{G} = \{1, 2, \dots, n\}$ and $g_\gamma(j) = j$.

Next, we define the γ -th verification matrix $A^{(V,\gamma)}$ in a way similar to how we defined the verification entries in the exactly k -sparse case. For each (i,j) , if the graph \mathcal{G} does not have an edge connecting i on the right to j on the left, then $a_{i,j}^{(V,\gamma)} = 0$. Otherwise, we set

$$a_{i,j}^{(V,\gamma)} = e^{i \theta_{ij}^{(V,\gamma)}},$$

where $\theta_{ij}^{(V,\gamma)}$ is drawn uniformly at random from $\{0, \pi/2(|\mathbf{G}| - 1), \pi/(|\mathbf{G}| - 1), 3\pi/2(|\mathbf{G}| - 1), \dots, \pi/2\}$.

Given an signal vector \mathbf{x} , signal noise \mathbf{z} , and measurement noise \mathbf{e} , the measurement operation produces a measurement vector $\mathbf{y} = A(\mathbf{x} + \mathbf{e})$. Since A can be partitioned into Γ identification and Γ verification rows, we think of the measurement vector \mathbf{y} as a collection of outcomes from Γ successive measurement operations such that

$$\mathbf{y}^{(I,\gamma)} = A^{(I,\gamma)}(\mathbf{x} + \mathbf{z}) + \mathbf{e}^{(I,\gamma)}$$

and

$$\mathbf{y}^{(V,\gamma)} = A^{(V,\gamma)}(\mathbf{x} + \mathbf{z}) + \mathbf{e}^{(V,\gamma)}$$

are the outcomes from the γ -th measurement and $\mathbf{y} = ((\mathbf{y}^{(I,\gamma)}, \mathbf{y}^{(V,\gamma)}) : 1 \leq \gamma \leq \Gamma)$.

C. Reconstruction for approximately k -sparse signals with noisy measurements

The decoding algorithm for this case extends the decoding algorithm presented earlier for the exactly k -sparse case by including the ideas presented in Section III-A. The total number of iterations for our algorithm are upper bounded by $4k$.

- 1) We initialize by setting the *signal estimate vector* $\hat{\mathbf{x}}(1)$ to the all-zeros vector 0^n , and for each $\gamma = 1, 2, \dots, \Gamma$, we set the *residual measurement identification/verification vectors* $\tilde{\mathbf{y}}^{(I,\gamma)}(1)$ and $\tilde{\mathbf{y}}^{(V,\gamma)}(1)$ to the decoder's observations $\mathbf{y}^{(I,\gamma)}$ and $\mathbf{y}^{(V,\gamma)}$.

Let $\mathcal{D}(1)$, the initial neighborly set, be the set of indices i for which, at which the magnitude corresponding to all verification and identification vectors is greater than δ/k , i.e.,

$$\mathcal{D}(1) = \bigcap_{\gamma=1}^{\Gamma} \left\{ i : |y_i^{(I,\gamma)}| > \frac{\delta}{k}, |y_i^{(V,\gamma)}| > \frac{\delta}{k} \right\}$$

vector $\mathbf{y}^{(V)}$, i.e., the set $\{i \leq m : \tilde{y}_i^{(V)}(1) \neq 0\}$. This step takes $\mathcal{O}(k)$ steps, since merely reading \mathbf{y} to check for the zero locations of $\mathbf{y}^{(V)}$ takes that long.

- 2) The t^{th} decoding iteration accepts as its input the t^{th} signal estimate vector $\hat{\mathbf{x}}^{(t)}$, the t^{th} neighbourly set $\mathcal{D}(t)$, and the t^{th} residual measurement identification/verification vectors $((\tilde{\mathbf{y}}^{(I,\gamma)}(t), \tilde{\mathbf{y}}^{(V,\gamma)}(t)) : \gamma = 1, 2, \dots, \Gamma)$. In $\mathcal{O}(1)$ steps it outputs the $(t+1)^{\text{th}}$ signal estimate vector $\hat{\mathbf{x}}^{(t+1)}$, the $(t+1)^{\text{th}}$ neighbourly set $\mathcal{D}(t+1)$, and the $(t+1)^{\text{th}}$ residual measurement identification/verification vectors $((\tilde{\mathbf{y}}^{(I,\gamma)}(t+1), \tilde{\mathbf{y}}^{(V,\gamma)}(t+1)) : \gamma = 1, 2, \dots, \Gamma)$ after the performing the following steps sequentially (each of which takes at most a constant number of atomic steps).
- 3) *Pick a random $i(t)$* : The decoder picks $i(t)$ uniformly at random from $\mathcal{D}(t)$
- 4) *Compute quantized phases*: For each $\gamma = 1, 2, \dots, \Gamma$, compute the *current identification angles*, $\hat{\theta}_t^{(I,\gamma)}$, and *current identification angles*, $\hat{\theta}_t^{(V,\gamma)}$ defined as follows:

$$\begin{aligned} \hat{\theta}_t^{(I,\gamma)} &= \left\lfloor \frac{2(|\mathbf{G}| - 1) \left(\angle y_{i(t)}^{(I,\gamma)} \pmod{\pi} \right)}{\pi} \right\rfloor \frac{\pi}{2(|\mathbf{G}| - 1)}, \\ \hat{\theta}_t^{(V,\gamma)} &= \left\lfloor \frac{2(|\mathbf{G}| - 1) \left(\angle y_{i(t)}^{(V,\gamma)} \pmod{\pi} \right)}{\pi} \right\rfloor \frac{\pi}{2(|\mathbf{G}| - 1)}. \end{aligned}$$

In the above, $\lfloor \cdot \rfloor$ denotes the closest integer function. Since there are $\Theta(n)$ different phase vectors, to perform this computation, $\mathcal{O}(\log n)$ precision and $\mathcal{O}(1)$ steps suffice.

For each $\gamma = 1, 2, \dots, \Gamma$, let $\hat{g}_\gamma^{(t)} = 2(|\mathbf{G}| - 1)\hat{\theta}_t^{(I,\gamma)}/\pi$ be the *current estimate of γ -th digit* and let $j(t)$ be the number whose representation in \mathbf{G} is $(\hat{g}_1^{(t)}, \hat{g}_2^{(t)}, \dots, \hat{g}_\Gamma^{(t)})$.

- 5) *Check if the current identification and verification angles correspond to a valid and unique j* : This step determines if $i(t)$ is a leaf node for $\mathcal{S}_\delta(\mathbf{x} - \hat{\mathbf{x}}(t))$. This operation is similar to the corresponding exact- k case. The main difference here is that we perform the verification operation on each of the Γ measurements separately and declare $i(t)$ as a leaf node only if it passes all the verification tests. The verification step for the γ -th measurement is given by the test:

$$\hat{\theta}_t^{(V,\gamma)} \stackrel{?}{=} \theta_{i(t),j(t)}^{(V,\gamma)}.$$

If the above test succeeds for every $\gamma = 1, 2, \dots, \Gamma$, we set $\Delta x(t)$ to $|\tilde{y}_{i(t)}^{(I,\gamma)}(t)|$ if $\angle y_{i(t)}^{(I,\gamma)} \in (-\pi/4, 3\pi/4]$, and $-|\tilde{y}_{i(t)}^{(I,\gamma)}(t)|$ if $\angle y_{i(t)}^{(I,\gamma)} \in (3\pi/4, 7\pi/4]$. Otherwise, we set $\Delta x(t) = 0$. This step requires at most Γ verification steps and therefore, can be completed in $\mathcal{O}(1)$ steps.

- 6) *Update $\hat{\mathbf{x}}(t+1)$, $\tilde{\mathbf{y}}(t+1)$, and $\mathcal{D}(t+1)$* : If the verification tests in the previous steps failed, there are no updates to be done, i.e., set $\hat{\mathbf{x}}(t+1) = \hat{\mathbf{x}}(t)$, $\tilde{\mathbf{y}}(t+1) = \tilde{\mathbf{y}}(t)$, and $\mathcal{D}(t+1) = \mathcal{D}(t)$. Otherwise, we first update the current signal estimate to $\hat{\mathbf{x}}(t+1)$ by setting the $j(t)$ -th coordinate to $\Delta x(t)$. Next, let i_1, i_2, i_3 be the possible neighbours of $j(t)$. We compute the *residual identification/verification vectors*

$\tilde{y}(t+1)$ at i_1, i_2, i_3 by subtracting the weight due to $\Delta x(t)$ at each of them. Finally, we update the neighbourly set by removing i_1, i_2 , and i_3 from $\mathcal{D}(t)$ to obtain $\mathcal{D}(t+1)$.

The decoding algorithm terminates after the T -th iteration, where $T = \min\{4k, \{t : \mathcal{D}(t+1) = \phi\}\}$.

IV. CONCLUSION

In this work we present SHO-FA – an algorithm for sparse recovery that requires an information-theoretically order-optimal number of measurements, bits over all measurements, and decoding time-complexity, and as a bonus, with non-zero probability it can handle “data-base queries”. The algorithm is robust to noisy signal tails and noisy measurements. The algorithm is “practical” (all constant factors involved are “small”), as validated by both our analysis, and simulations. Our algorithm can reconstruct signals that are sparse in any basis that is known *a priori* to both the encoder and decoder, and works for a large ensemble of “sparse” measurement matrices A .

As future work we aim to use recent work by [32] and [34] to improve respectively the constant factors in our performance analysis and reconstruction error.

REFERENCES

- [1] E. J. Candès, J. K. Romberg, and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, pp. 489–509, 2006.
- [2] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, pp. 1289–1306, 2006.
- [3] E. J. Candès, “The restricted isometry property and its implications for compressed sensing,” *Comptes Rendus Mathématique*, vol. 346, no. 9-10, pp. 589–592, 2008.
- [4] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, “A simple proof of the restricted isometry property for random matrices,” *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, December 2008.
- [5] A. Cohen, R. DeVore, and W. Dahmen, “Compressed sensing and best k-term approximation,” *Journal of the AMS*, vol. 22, pp. 211–231, 2009.
- [6] A. C. Gilbert, Y. Li, E. Porat, and M. J. Strauss, “Approximate sparse recovery: optimizing time and measurements,” in *Proceedings of the 42nd ACM symposium on Theory of computing*, ser. STOC ’10, 2010, pp. 475–484.
- [7] E. Price and D. P. Woodruff, “(1 + ϵ)-approximate sparse recovery,” in *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, ser. FOCS ’11, 2011, pp. 295–304.
- [8] J. A. Tropp and A. C. Gilbert, “Signal recovery from random measurements via orthogonal matching pursuit,” *IEEE Transactions on Information Theory*, pp. 4655–4666, 2007.
- [9] D. L. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck, “Sparse solution of underdetermined systems of linear equations by stagewise orthogonal matching pursuit,” *IEEE Transactions on Information Theory*, pp. 1094–1121, 2012.
- [10] R. Berinde, P. Indyk, and M. Ruzic, “Practical near-optimal sparse recovery in the ℓ_1 norm,” *Proceedings of the Annual Allerton conference*, 2008.
- [11] R. Berinde and P. Indyk, “Sequential sparse matching pursuit,” *Proceedings of the Annual Allerton conference*, 2009.
- [12] A. Gilbert and P. Indyk, “Sparse recovery using sparse matrices,” *Proceedings of IEEE*, vol. 98, no. 6, pp. 937–947, 2010.
- [13] K. D. Ba, P. Indyk, E. Price, and D. Woodruff, “Lower bounds for sparse recovery,” *Proceedings of the Symposium on Discrete Algorithms*, 2010.
- [14] I. S. Reed and G. Solomon, “Polynomial codes over certain finite fields,” *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300 – 304, Jun 1960.
- [15] F. Parvaresh and B. Hassibi, “Explicit measurements with almost optimal thresholds for compressed sensing,” in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2008, pp. 3853–3856.
- [16] S. Kudekar and H. Pfister, “The effect of spatial coupling on compressive sensing,” in *Proc. Annual Allerton Conf. on Commun., Control, and Comp*, 2010.
- [17] M. Mitzenmacher and G. Varghese, “Biff (bloom filter) codes: Fast error correction for large data sets,” *To appear in ISIT*, 2012.
- [18] S. Jafarpour, W. Wu, B. Hassibi, and R. Calderbank, “Efficient and robust compressed sensing using optimized expander graphs,” *IEEE Transactions on Information Theory*, vol. 55, no. 9, pp. 4299–4308, 2009.
- [19] M. Akçakaya and V. Tarokh, “Shannon-theoretic limits on noisy compressive sampling,” *IEEE Transactions on Information Theory*, vol. 56, no. 1, pp. 492–504, Jan. 2010.
- [20] Y. Wu and S. Verdú, “Rényi information dimension: Fundamental limits of almost lossless analog compression,” *IEEE Transaction on Information Theory*, vol. 56, no. 8, pp. 3721–3748, 2010.
- [21] Y. Wu and S. Verdú, “Optimal phase transitions in compressed sensing,” *ArXiv.org e-Print archive*, arXiv:1111.6822 [cs.IT], 2011.
- [22] D. L. Donoho, A. Javanmard, and A. Montanari, “Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing,” *ArXiv.org e-Print archive*, arXiv:1112.0708 [cs.IT], 2011.
- [23] A. K. Fletcher, S. Rangan, and V. K. Goyal, “Necessary and sufficient conditions for sparsity pattern recovery,” *IEEE Transactions on Information Theory*, vol. 55, no. 12, pp. 5758 – 5772, Dec. 2009.
- [24] M. J. Wainwright, “Information-theoretic limitations on sparsity recovery in the high-dimensional and noisy setting,” *IEEE Transactions on Information Theory*, vol. 55, no. 12, pp. 5728 – 5741, Dec. 2009.
- [25] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, and A. Kabbani, “Counter braids: a novel counter architecture for per-flow measurement,” in *Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, ser. SIGMETRICS ’08, 2008, pp. 121–132.

- [26] Y. Lu, A. Montanari, and B. Prabhakar, “Counter braids: Asymptotic optimality of the message passing decoding algorithm,” in *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, Sept. 2008, pp. 209–216.
- [27] Y. Plan and R. Vershynin, “One-bit compressed sensing by linear programming,” *ArXiv.org e-Print archive*, arXiv:1109.4299 [cs.IT], 2011.
- [28] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, “Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors,” *ArXiv.org e-Print archive*, arXiv:1104.3160 [cs.IT], 2011.
- [29] A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin, “Algorithmic linear dimension reduction in the ℓ_1 norm for sparse vectors,” in *Allerton 2006 (44th Annual Allerton Conference on Communication, Control, and Computing)*, 2006.
- [30] G. Cormode and S. Muthukrishnan, “Combinatorial algorithms for compressed sensing,” in *Information Sciences and Systems, 2006 40th Annual Conference on*, March 2006, pp. 198–201.
- [31] S. Sarvotham, D. Baron, and R. Baraniuk, “Sudocodes - fast measurement and reconstruction of sparse signals,” in *Information Theory, 2006 IEEE International Symposium on*, July 2006, pp. 2804–2808.
- [32] M. T. Goodrich and M. Mitzenmacher, “Invertible bloom lookup tables,” *ArXiv.org e-Print archive*, arXiv:1101.2245 [cs.DS], 2011.
- [33] D. A. Spielman, “Linear-time encodable and decodable error-correcting codes,” in *STOC’95*, 1995, pp. 388–397.
- [34] E. Price, “Efficient sketches for the set query problem,” in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA ’11, 2011, pp. 41–56.
- [35] S. Pawar and K. Ramchandran, personal communication, 2012.
- [36] I. G. Shevtsova, “An improvement of convergence rate estimates in the lyapunov theorem,” *Doklady Mathematics*, vol. 82, no. 3, pp. 862–864, 2010.
- [37] S. N. Bernstein, “On certain modifications of chebyshev’s inequality,” *Doklady Akademii Nauk SSSR*, vol. 17, no. 6, pp. 275–277, 1937.
- [38] D. Guo, J. Luo, L. Zhang, and K. Shen, “Compressed neighbor discovery for wireless networks,” *CoRR*, vol. abs/1012.1007, 2010.
- [39] M. Alekhovich, “Linear diophantine equations over polynomials and soft decoding of reed-solomon codes,” *IEEE Transactions on Information Theory*, pp. 2257–2265, 2005.
- [40] R. C. Singleton, “Maximum distance q -nary codes,” *IEEE Transactions on Information Theory*, vol. 10, no. 2, pp. 116–118, 1964.
- [41] P. Indyk and M. Ruzic, “Near-optimal sparse recovery in the ℓ_1 norm,” in *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, ser. FOCS ’08, 2008, pp. 199–207.
- [42] M. Akçakaya and V. Tarokh, “A frame construction and a universal distortion bound for sparse representations,” *IEEE Transactions on Signal Processing*, pp. 2443–2450, 2008.
- [43] S. Aeron, V. Saligrama, and M. Zhao, “Information theoretic bounds for compressed sensing,” *Information Theory, IEEE Transactions on*, vol. 56, no. 10, pp. 5111–5130, Oct. 2010.
- [44] R. Roth, *Introduction to coding theory*. Cambridge University Press, 2006.
- [45] S. Hoory, N. Linial, and A. Wigderson, “Expander graphs and their applications,” *Bulletin (New series) of the American Mathematical Society*, vol. 43, pp. 439–561, 2006.
- [46] H. Y. Cheung, T. C. Kwok, and L. C. Lau, “Fast matrix rank algorithms and applications,” *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012.
- [47] D. J. C. Mackay, *Information theory, inference and learning algorithms*, first edition ed. Cambridge University Press, Jun.
- [48] J. T. Schwartz, “Fast probabilistic algorithms for verification of polynomial identities,” *Journal of the ACM*, vol. 27, no. 4, p. 701–717, 1980.
- [49] R. Zippel, “Probabilistic algorithms for sparse polynomials,” *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pp. 216–226, 1979.
- [50] M. Karonński and T. Łuczak, “The phase transition in a random hypergraph,” *J. Comput. Appl. Math.*, vol. 142, no. 1, pp. 125–135, May 2002.

APPENDIX

A. Proof of Lemma 1

Proof: It suffices to prove the desired property for all $\mathcal{S}(\mathbf{x})$ of size exactly k . Let $\mathcal{S}'(\mathbf{x}) \subseteq \mathcal{S}(\mathbf{x})$. Let $\{(s_1, t_1), (s_2, t_2), \dots, (s_{d|\mathcal{S}'(\mathbf{x})|}, t_{d|\mathcal{S}'(\mathbf{x})|})\}$ be the set of outgoing edges from $\mathcal{S}'(\mathbf{x})$. Without loss of generality, we assume these edges are drawn in the following manner.

In the initialization stage, we “split” every node on the right of \mathcal{G} to $dn/c k$ “virtual” nodes²². Each virtual node represents a “true” node on the right. We maintain a set of “remaining” virtual nodes, which we will select and remove virtual nodes from.

To draw the edges, we visit the nodes in $\mathcal{S}'(\mathbf{x})$ (on the left of \mathcal{G}) sequentially. For each node, we select uniformly at random a set of d distinct virtual nodes from the remaining virtual node set. We form d edges by connecting this node in $\mathcal{S}'(\mathbf{x})$ and the true nodes on the right that those d selected virtual nodes represent. After the d edges are formed, we remove the d selected virtual nodes from the remaining virtual node set, and proceed to the next node in $\mathcal{S}'(\mathbf{x})$.

In this way, we generate a bipartite graph that is both d left-regular and $dn/c k$ right-regular; that is, each node on the left has a degree of d and each node on the right has a degree of $dn/c k$. By using standard arguments of

²²We assume $dn/c k$ is an integer, with the understanding that in practice one can always increase c to make $dn/c k$ integer while the “fail-to-expand” probability is still bounded by the desired target ϵ .

sequential implementation of random experiments, one can verify that the graph generated in this way is chosen uniformly at random from all bipartite graphs that are both d left-regular and dn/ck right-regular.

For each $i = 1, 2, \dots, d|\mathcal{S}'(\mathbf{x})|$, the probability that the edge (s_i, t_i) reaches an “old” true node (on the right) that is already reached by those edges generated ahead of (s_i, t_i) is upper bounded as

$$\begin{aligned} \Pr_{\mathcal{G}}(t_i \in \{t_1, \dots, t_{i-1}\}) &\leq \frac{(i-1)}{ck} \\ &\leq \frac{d|\mathcal{S}'(\mathbf{x})|}{ck}. \end{aligned}$$

Let $N(\mathcal{S}'(\mathbf{x}))$ be the set of all neighboring nodes of the nodes in $\mathcal{S}'(\mathbf{x})$. The size of $N(\mathcal{S}'(\mathbf{x}))$ is no more than $2d|\mathcal{S}'(\mathbf{x})|/3$ if and only if out of $d|\mathcal{S}'(\mathbf{x})|$ edges, there exists a set of at least $d|\mathcal{S}'(\mathbf{x})|/3$ edges fail to reach “new” nodes (on the right). Exploiting this observation, we have

$$\begin{aligned} &\Pr_{\mathcal{G}}(|N(\mathcal{S}'(\mathbf{x}))| \leq 2d|\mathcal{S}'(\mathbf{x})|/3) \\ &= \Pr_{\mathcal{G}}\left(\bigcup_{\substack{\sigma \subseteq \{1, \dots, d|\mathcal{S}'(\mathbf{x})|\} \\ |\sigma| \geq d|\mathcal{S}'(\mathbf{x})|/3}} \bigcap_{i \in \sigma} \{t_i \in \{t_1, \dots, t_{i-1}\}\}\right) \\ &= \Pr_{\mathcal{G}}\left(\bigcup_{\substack{\sigma \subseteq \{1, \dots, d|\mathcal{S}'(\mathbf{x})|\} \\ |\sigma| = d|\mathcal{S}'(\mathbf{x})|/3}} \bigcup_{\sigma' \supseteq \sigma} \bigcap_{i \in \sigma'} \{t_i \in \{t_1, \dots, t_{i-1}\}\}\right) \\ &\leq \binom{d|\mathcal{S}'(\mathbf{x})|}{d|\mathcal{S}'(\mathbf{x})|/3} \left(\frac{d|\mathcal{S}'(\mathbf{x})|}{ck}\right)^{d|\mathcal{S}'(\mathbf{x})|/3}. \end{aligned}$$

Consequently, the probability that there exists one $\mathcal{S}'(\mathbf{x}) \subseteq \mathcal{S}(\mathbf{x})$ so that $|N(\mathcal{S}'(\mathbf{x}))| \leq 2d|\mathcal{S}'(\mathbf{x})|/3$ can be bounded by

$$\begin{aligned} &\Pr_{\mathcal{G}}(\cup_{\mathcal{S}'(\mathbf{x}) \subseteq \mathcal{S}(\mathbf{x})} \{|N(\mathcal{S}'(\mathbf{x}))| \leq 2|\mathcal{S}'(\mathbf{x})|\}) \\ &\leq \sum_{\mathcal{S}'(\mathbf{x}) \subseteq \mathcal{S}(\mathbf{x})} \binom{d|\mathcal{S}'(\mathbf{x})|}{d|\mathcal{S}'(\mathbf{x})|/3} \left(\frac{d|\mathcal{S}'(\mathbf{x})|}{ck}\right)^{d|\mathcal{S}'(\mathbf{x})|/3} \\ &= \sum_{j=1}^k \binom{k}{j} \left(\frac{dj}{dj/3}\right) \left(\frac{dj}{ck}\right)^{dj/3} \\ &\leq \sum_{j=1}^k \left(\frac{ke}{j}\right)^j (3e)^{dj/3} \left(\frac{dj}{ck}\right)^{dj/3} \tag{7} \end{aligned}$$

$$\begin{aligned} &= \sum_{j=1}^k \left(\frac{ke}{j} \left(\frac{3dje}{ck}\right)^{d/3}\right)^j \\ &\leq \sum_{j=1}^{\lceil \sqrt{k} \rceil} \left(\frac{ke}{j} \left(\frac{3dje}{ck}\right)^{d/3}\right)^j + \sum_{j=\lfloor \sqrt{k} \rfloor}^k \left(\frac{ke}{j} \left(\frac{3dje}{ck}\right)^{d/3}\right)^j \\ &\leq \sqrt{k} \left(\frac{ke}{\sqrt{k}} \left(\frac{3d\sqrt{k}e}{ck}\right)^{d/3}\right) + \sum_{j=\lfloor \sqrt{k} \rfloor}^{\infty} \left(e \left(\frac{3de}{c}\right)^{d/3}\right)^j \tag{8} \end{aligned}$$

$$\leq \left(\frac{3de}{c}\right)^{d/3} k^{-d/6} e + \exp(-\theta(\sqrt{k})) \tag{9}$$

$$= \mathcal{O}(k^{-d/6}). \tag{10}$$

In the above, the inequality in (7) follows from Stirling's approximation; the upper bound in (8) is derived by noting that the first term in the sum takes its maximum when $j = \lfloor \sqrt{k} \rfloor$ and the second term is maximum when $j = k$; (9) is obtained by noting that the second term is a geometric progression.

Finally, we plug in the choice of $d = 7$ to complete the proof. \blacksquare

B. Proof of Lemma 2

Suppose each set of size k of $\mathcal{S}(\mathbf{x})$ nodes on the left of \mathcal{G} has strictly more than $d/2$ times as many nodes neighbouring those in $\mathcal{S}(\mathbf{x})$, as there are in $\mathcal{S}(\mathbf{x})$. Then by standard arguments in the construction of expander codes [33], this implies the existence of a linear code of rate at least $1 - m/n$, and with relative minimum distance at least k/n .²³ But by the Hamming bound [44], it is known that codes of minimum distance δ can have rate at most $1 - H(\delta)$, where $H(\cdot)$ denotes the binary entropy function. Since $k = (n)$, $\delta = k/n \rightarrow 0$. But in this regime $1 - H(\delta) \rightarrow 1 - \delta \log(1/\delta)$. Comparing $(k/n) \log(n/k)$ with m/n gives the required result. \blacksquare

C. Proof of Lemma 3

For any set of nodes S in the graph \mathcal{G} , we define $N(S)$ as the set of neighboring nodes of the nodes in S . For any set $\mathcal{S}'(\mathbf{x}) \subseteq \mathcal{S}(\mathbf{x})$, we define β as the portion of the nodes in $N(\mathcal{S}'(\mathbf{x}))$ that are $\mathcal{S}'(\mathbf{x})$ -leaf nodes.

First, each node $v \in N(\mathcal{S}'(\mathbf{x}))$ is of one of the following two types:

- 1) It has only one neighboring node in $\mathcal{S}'(\mathbf{x})$, on the left of \mathcal{G} . By the definition of β , the number of nodes in $N(\mathcal{S}'(\mathbf{x}))$ of this type is $\beta|N(\mathcal{S}'(\mathbf{x}))|$.
- 2) It has at least two neighboring nodes in $\mathcal{S}'(\mathbf{x})$, on the left of \mathcal{G} . The number of nodes in $N(\mathcal{S}'(\mathbf{x}))$ of this type is $(1 - \beta)|N(\mathcal{S}'(\mathbf{x}))|$.

We have two observations. First, since the degree of each node in $\mathcal{S}'(\mathbf{x})$ is d , the total number of edges from $\mathcal{S}'(\mathbf{x})$ to $N(\mathcal{S}'(\mathbf{x}))$ is at most $d|\mathcal{S}'(\mathbf{x})|$ and the number of nodes in $N(\mathcal{S}'(\mathbf{x}))$ is at most $d|\mathcal{S}'(\mathbf{x})|$.

Second, the total number of edges entering $N(\mathcal{S}'(\mathbf{x}))$ from $\mathcal{S}'(\mathbf{x})$ is at least

$$\beta|N(\mathcal{S}'(\mathbf{x}))| + 2(1 - \beta)|N(\mathcal{S}'(\mathbf{x}))| = (2 - \beta)|N(\mathcal{S}'(\mathbf{x}))|,$$

as the number of neighboring nodes for the nodes of Type 1 is one and of Type 2 is at least two.

Combining the above two observations, we can get the following inequality:

$$(2 - \beta)d|N(\mathcal{S}'(\mathbf{x}))|/3 \leq d|\mathcal{S}'(\mathbf{x})|.$$

According to the setting of the Lemma, we also have $|N(\mathcal{S}'(\mathbf{x}))| \geq 2d/|\mathcal{S}'(\mathbf{x})|3$. Therefore, it follows that

$$2(2 - \beta)d|\mathcal{S}'(\mathbf{x})|/3 \leq d|\mathcal{S}'(\mathbf{x})|,$$

and consequently $\beta \geq 1/2$. \blacksquare

D. Proof of Lemma 4

Consider the algorithm \mathcal{A} that proceeds as follows. First, among the set of all right nodes that neighbour j , check if there exists a node i such that $y_i^{(I)} = y_i^{(V)} = 0$. If there exists such a node, then output $\hat{x}_j = 0$. Otherwise, check if there exists a $\mathcal{S}(\mathbf{x})$ -leaf node among the neighbours of j . This check can be performed by using verification and identification observations as described for the SHO-FA reconstruction algorithm. If there exists a leaf node, say i , then output $\hat{x}_j = |y_i|$. Else, the algorithm terminates without producing any output.

Two see that the above algorithm satisfies the claimed properties, consider the following two cases.

Case 1: $x_j = 0$. In this case, $\hat{x}_j = 0$ is output if at least one neighbour of j lies outside $N(\mathcal{S}(\mathbf{x}))$. Since $N(\mathcal{S}(\mathbf{x}))$

²³For the sake of completeness we sketch such an argument here. Given such an expander graph \mathcal{G} , one can construct a $n \times k$ binary matrix A with 1s in precisely those (i, j) locations where the i th node on the left is connected with the j th node on the right. Treating this matrix A as the parity check matrix of a code over a block-length n implies that the rate of the code is at least k/n , since the parity-check matrix imposes at most k constraints on the n bits of the codewords. Also, the minimum distance is at least k . Suppose not, i.e. there exists a codeword in this linear code of weight less than k . Let the support of this codeword be denoted $\mathcal{S}(\mathbf{x})$. Then by the expansion property of \mathcal{G} , there are strictly more than $|\mathcal{S}(\mathbf{x})|d/2$ neighbours of $\mathcal{S}(\mathbf{x})$. But this implies that there is at least one node, say v , neighboring $\mathcal{S}(\mathbf{x})$ which has exactly one neighbor in $\mathcal{S}(\mathbf{x})$. But then the constraint corresponding to v cannot be satisfied, leading to a contradiction.

has at most dk elements, the probability that a neighbour of j lies inside $N(\mathcal{S}(\mathbf{x}))$ is at most $dk/c = d/c$. Thus, the probability that none of the neighbours of j lie outside $N(\mathcal{S}(\mathbf{x}))$ is at least $(1 - (d/c)^d)$. The algorithm incorrectly reconstructs x_j if all neighbours of j lie within $N(\mathcal{S}(\mathbf{x}))$ and SHO-FA incorrectly identifies one of these nodes as a leaf node. By the analysis of SHO-FA, this event occurs with probability $o(1/k)$.

Case 2: $x_j \neq 0$. For \mathcal{A} to produce the correct output, it has to identify one of the neighbours of j as a leaf. The probability that there exists a leaf among the neighbours of j is at least $(1 - (d/c)^d)$ by an argument similar to the previous case. Similarly, the probability of erroneous identification is $o(1/k)$. ■

E. Proof of Lemma 5

First, we find an upper bound on the maximum possible phase displacement in y_i due to fixed noise vectors \mathbf{z} and \mathbf{e} . Let $\Delta\theta_i$ be the difference in phase between the "noiseless" output $(A'\mathbf{x})_i$ and the actual output $y_i = (A'(\mathbf{x} + \mathbf{z}) + \mathbf{e})_i$. Figure 11a shows this geometrically. By a straightforward geometric argument, for fixed \mathbf{z} and \mathbf{e} , the phase displacement $\Delta\theta_i$ is upper bounded by $\pi|(A'\mathbf{z})_i + e_i|/|(A'\mathbf{x})_i|$. Since i is a leaf node for $\mathcal{S}(\mathbf{x})$, $|(A'\mathbf{x})_i| \geq |\delta/k|$. Therefore,

$$\Delta\theta_i \leq \pi|(A'\mathbf{z})_i + e_i|k/\delta.$$

Thus,

$$\begin{aligned} \Pr_{\mathbf{z}, \mathbf{e}}(\Delta\theta_i > \alpha) &\leq \Pr_{\mathbf{z}, \mathbf{e}}\left(|(A'\mathbf{z})_i + e_i|k/\delta > \alpha\right) \\ &= \Pr_{\mathbf{z}, \mathbf{e}}\left(|(A'\mathbf{z})_i + e_i| > \alpha\delta/\pi k\right) \end{aligned}$$

Since each z_j is a Gaussian with zero mean and variance σ_z^2 , $(A'\mathbf{z})_i$ is a Complex Gaussian with zero mean and variance at most $n\sigma_z^2$. Further, each row of A' has at most dn/c non-zero entries. Therefore, $(A'\mathbf{z})_i + e_i$ is a zero mean complex Gaussian with variance at most $(dn/c)\sigma_z^2 + \sigma_e^2$.

The expected value of $\Delta\theta_i$ is bounded as follows:

$$\begin{aligned} E_{\mathbf{z}, \mathbf{e}}(\Delta\theta_i) &\leq E_{\mathbf{z}, \mathbf{e}}(\pi|(A'\mathbf{z})_i + e_i|k/\delta) \\ &\leq \frac{\pi k}{\delta} \int_0^\infty \sqrt{\frac{2}{\pi(dn\sigma_z^2/c + \sigma_e^2)}} l e^{-l^2/2(dn\sigma_z^2/c + \sigma_e^2)} dl \\ &= \sqrt{\frac{2\pi k^2(dn\sigma_z^2/c + \sigma_e^2)}{\delta^2}}. \end{aligned}$$

Finally, applying standard bounds on the tail probabilities of Gaussian random variables, the required probability is upper bounded by $e^{-(\alpha^2/2\pi)}/2$. ■

F. Probability of error

An error occurs only if one of the following take place:

- 1) The underlying graph \mathcal{G} is not an $\mathcal{S}(\mathbf{x})$ -expander. This probability can be made $o(1/k)$ by choosing $m = ck$, where the constant c is determined by Lemma 1
- 2) The phase noise in $\tilde{y}_{i(t)}(t)$ leads to an incorrect decoding of $\hat{\theta}_t^{(I, \gamma)}$ or $\hat{\theta}_t^{(V, \gamma)}$ for some γ and t .

Note that the phase noise in $\tilde{y}_{i(t)}(t)$ consists:

- a) The contribution due to noise vectors \mathbf{z} and \mathbf{e} , and
- b) The contribution due to the noise propagated while computing each $\tilde{y}_{i(t)}(\tau)$ from $\tilde{y}_{i(t)}(\tau - 1)$ for $\tau \leq t$.

The contribution due to the first term is bounded by Lemma 5. Thus, for a target error probability ϵ' , we choose $\alpha = (1/2) \log 1/2\epsilon'$, giving a contribution to the phase noise of at most

$$\frac{\log(1/2\epsilon')}{2} \sqrt{\frac{2\pi k^2(dn\sigma_z^2/c + \sigma_e^2)}{\delta^2}}.$$

To bound the contribution due to the second term, note that at each iteration t , any error in reconstruction of $\hat{x}_{j(t)}$ potentially adds to reconstruction error in all future iterations t' for which there is a path from $j(t)$ to $j(t')$. Since the restriction of \mathcal{G} to $\mathcal{S}(\mathbf{x})$ and its neighbours is a sparse graph, it follows from [50] that, with probability $1 - o(1/k)$, it consists only of disjoint components of size $\mathcal{O}(\log k)$ (see [34] for such an analysis). Thus, the magnitude error in reconstruction of $\hat{x}_{j(t)}$ due to noisy reconstructions in previous iterations is

$$\mathcal{O}\left((\log k)^2 \log(1/\epsilon') \sqrt{(n\sigma_z^2/k + \sigma_e^2)}\right). \quad (11)$$

Thus, the phase displacement in each $y_i^{(I,\gamma)}$ and $y_i^{(V,\gamma)}$ is

$$\mathcal{O}\left((\log k)^2 \log(1/\epsilon') \sqrt{\frac{k^2(n\sigma_z^2/k + \sigma_e^2)}{\delta^2}}\right).$$

Therefore, as long as

$$(\log k)^2 \log(1/\epsilon') \sqrt{\frac{k^2(n\sigma_z^2/k + \sigma_e^2)}{\delta^2}} = o\left(n^{-1/\Gamma}\right), \quad (12)$$

the probability of any single phase being incorrectly detected is upper bounded by ϵ' . Since we there are a total of $8\Gamma k$ possible phase measurements, we choose $\epsilon' = \mathcal{O}(1/\Gamma k^2)$ to achieve a target error probability $\mathcal{O}(1/k)$.

- 3) The verification step passes for each measurement in the t -th measurement, even though $i(t)$ is not a leaf node for $\mathcal{S}_\delta^c(\mathbf{x})$.
- 4) $\mathcal{D}(T) \neq A'$, i.e., the algorithm terminates without recovering all x_j 's. Note that similar to the exact k -sparse case, in each iteration t , by Lemma 3, the probability that $i(t)$ is a leaf node for $\mathcal{S}_\delta(\mathbf{x} - \hat{\mathbf{x}}(t))$ at least $1/2$. However, due to noise, there is a non-zero probability that even when $i(t)$ is a leaf node, it does not pass the verification tests. We know from the analysis for the previous case that this probability is $\mathcal{O}(1/k)$ for each $i(t)$. Therefore, the probability that a randomly picked $i(t)$ passes the verification test is $1/2 - \mathcal{O}(1/k)$. Thus, in expectation, the number of iterations required by the algorithm is $2k/(1 - \mathcal{O}(1/k))$. By concentration arguments, it follows that the probability that the algorithm does not terminate in $4k$ iterations is $o(1/k)$ as k grows without bound.

G. Estimation error

Next, we bound the error in estimating $\hat{\mathbf{x}}$. We first find an upper bound on $\|\hat{\mathbf{x}} - \mathbf{x}_{\mathcal{S}_\delta^c}\|_1$ that holds with a high probability. Applying the bound in (11), for each $t = 1, 2, \dots, T$,

$$|x_{j(t)} - \hat{x}_{j(t)}| = \mathcal{O}\left((\log k)^2 \log(1/\epsilon') \sqrt{(n\sigma_z^2/k + \sigma_e^2)}\right)$$

with probability $1 - o(1/k)$. Therefore,

$$\begin{aligned} \|\hat{\mathbf{x}} - \mathbf{x}_{\mathcal{S}_\delta^c}\|_1 &= \sum_{\substack{1 \leq t \leq T \\ t: j(t) \notin \mathcal{S}_\delta}} |\hat{x}_j - x_j| + \sum_{\substack{1 \leq t \leq T \\ t: j(t) \in \mathcal{S}_\delta}} |\hat{x}_j| \\ &\leq \sum_{j \notin \mathcal{S}_\delta} |\hat{x}_j - x_j| + \sum_{j \in \mathcal{S}_\delta} |\hat{x}_j - x_j| + \sum_{j \in \mathcal{S}_\delta} |x_j| \\ &= \mathcal{O}\left(k(\log k)^2 \log(1/\epsilon') \sqrt{(n\sigma_z^2/k + \sigma_e^2)}\right) + \delta \text{ w.p. } 1 - o(1/k) \\ &= \mathcal{O}\left(k(\log k)^2 \log(1/\epsilon') \left(\sqrt{n\sigma_z^2/k} + \sigma_e\right) + \delta\right) \end{aligned} \quad (13)$$

Next, note that $\|\mathbf{z}\|_1 = \sum_{j=1}^n |z_j|$ and $\|\mathbf{e}\|_1 = \sum_{i=1}^m |e_i|$. Since each z_j is a Gaussian random variable with variance σ_z^2 , The expected value of $|z_j|$ is $\sigma_z \sqrt{2/\pi}$. Therefore, for every $\epsilon' > 0$, for n large enough,

$$\Pr(\|\mathbf{z}\|_1 < (1/2)n\sigma_z \sqrt{2/\pi}) < \epsilon'. \quad (14)$$

Similarly, for m large enough,

$$\Pr(\|\mathbf{e}\|_1 < (1/2)ck\sigma_e \sqrt{2/\pi}) < \epsilon'. \quad (15)$$

Combining inequalities (13)- (15), we have, with a high probability,

$$\begin{aligned} \|\hat{\mathbf{x}} - \mathbf{x}_{\mathcal{S}_\delta}\|_1 &= \mathcal{O}\left(k(\log k)^2 \log(1/\epsilon') \left(\frac{\|\mathbf{z}\|_1}{\sqrt{nk}} + \frac{\|\mathbf{e}\|_1}{k}\right)\right) + \delta \\ &= \mathcal{O}\left(\sqrt{\frac{k}{n}}(\log k)^2 \log(1/\epsilon') \|\mathbf{z}\|_1 + (\log k)^2 \log(1/\epsilon') \|\mathbf{e}\|_1\right) + \delta. \end{aligned} \quad (16)$$

Next, applying the bound in (3), we obtain

$$\|\hat{\mathbf{x}} - \mathbf{x}\|_1 \leq \mathcal{O}\left(\sqrt{\frac{k}{n}}(\log k)^2 \log(1/\epsilon') \|\mathbf{z}\|_1 + (\log k)^2 \log(1/\epsilon') \|\mathbf{e}\|_1\right) + 2\delta \quad (17)$$

with a high probability.

H. Proof of Theorem 2

Finally, to complete the proof of Theorem 2, we let $\delta = \min\{\mathcal{O}(n\sigma_z), o(1)\}$. By (14) with a high probability, $\delta = \mathcal{O}(\|\mathbf{z}\|)$. Finally, recall the assumption that $k = \mathcal{O}(n^{1-\Delta})$. Applying these to the bound obtained in (17), we get

$$\|\hat{\mathbf{x}} - \mathbf{x}\|_1 \leq C (\|\mathbf{z}\|_1 + (\log k)^2 \|\mathbf{e}\|_1)$$

for an appropriate constant $C = C(\epsilon)$.

I. Simulation Results

This section describes simulations that use synthetic data. The k -sparse signals used here are generated by randomly choosing k locations for non-zero values and setting the non-zero values to 1. The contours in each plot show the probability of successful reconstruction (the lighter the color, the higher the probability of reconstruction). The probability of error at each data point in the plots was obtained by running multiple simulations (400 in Fig 13 and Fig 14, and 200 in Fig 15) and noting the fraction of simulations which resulted in successful reconstruction.




Figure 13. **Exactly sparse signal and noiseless measurements – reconstruction performance for fixed signal length n :** The y -axis denotes the number of measurements m , and the x -axis denotes the sparsity k , for fixed signal length $n = 1000$. The simulation results show that the number of measurements m grows roughly proportional to the sparsity k for a fixed probability of reconstruction error. Also note that there is a sharp transition in reconstruction performance once the number of measurements exceeds a linear multiple of k . The red line denotes the curve where the probability of successful reconstruction equals 0.98. For $k = 150$, the probability of success equals 0.98 when $m = 450$ and $c = m/k = 3$.




Figure 14. **Exactly sparse signal and noiseless measurements – reconstruction performance for fixed sparsity k :** The number of measurements m are plotted on the y -axis, plotted against $\log(n)$ on the x -axis – the sparsity k is fixed to be 20. Note that there is *no* scaling of m with n , as guaranteed by our theoretical bounds.




Figure 15. **Approximately sparse signal and noisy measurements – reconstruction performance for fixed signal-length n :** As in Fig 13, the y -axis denotes the number of measurements m , and the x -axis denotes the sparsity k , for fixed signal length $n = 1000$. In this case, we set $\sigma_z = 0.03$, and allowed relative reconstruction error of at most 0.3.